



Employment and
Social Development Canada

Emploi et
Développement social Canada

Canada

Employment & Social Development Canada Artificial Intelligence Strategy

Version 3.0

Status: Draft

This document presents the Department of Employment and Social Development artificial intelligence strategy developed by the Chief Data office. The strategy details how we will strategically orient ourselves to obtain maximum value for Canadians in our AI investments and sets the stage for AI Policy and Governance to ensure it's used responsibly.

Table of Contents

Introductory Remarks	3
Strategic Initiatives	3
1. Why an AI Strategy?	4
1.2 What is Artificial Intelligence?	4
1.3 AI Strategy Objectives	5
2. Grounded in Experience	8
2.2 Our Current and Future Priorities for AI Use Cases	8
3. Demystifying Artificial Intelligence	10
3.1 Communicating ESDC's AI Strategy	10
3.2 Background on the Artificial Intelligence Definition	10
3.3a Text Classification	12
3.3b Information Retrieval	14
3.3c Question Answering (for Chat-bots)	15
3.3d Text Generation / Document Creation	18
3.3e Computer Vision	20
3.3f Client Segmentation	22
3.3g Strategic Optimization	24
4. AI Governance & Policy at ESDC	26
4.2 Proposed AI Governance	28
Preparing our existing committees for Artificial Intelligence initiatives	31
4.3 AI Policy: Key Dimensions of Consideration	32
4.3a Legal	32
4.3b Privacy	34
4.3c Transparency	36
4.3d Integrity & Security	38
4.3e Bias	39
4.3f Impact on the Workforce	41
4.3g Performance Measurement	42
4.3h Value	43
4.3i Other Policy Considerations	43
5. The Artificial Intelligence Value Proposition	45
5.2 Valuing ESDC's Data	45
ESDC's Data Strategy	46

ESDC AI Strategy

DRAFT - V3.0

5.3 ESDC's Analytics Program: Building for the Future.....	48
Analytics Program Oversight	49
5.4 The Need for Internal AI Capacity at ESDC.....	51
ESDC's Machine Learning Seminar	51
5.5 Obtaining Maximum Value with Vendors	52
<i>Vendor Guiding Principles</i>	52
6. Links to Other Departmental Initiatives	55

Introductory Remarks

Welcome to the Employment and Social Development Artificial Intelligence Strategy Landing Page

Artificial Intelligence is changing the world as we know it, and our clients' expectations and aspirations are changing along with it. This interactive AI Strategy website has been created to highlight how ESDC is approaching AI, both now and into the future.

Strategic Initiatives

1. Develop a modern AI suite to transform the way ESDC delivers service to Canadians
2. Engage across the organization to promote AI and coordinate initiatives
3. Develop a policy for acceptable AI use in light of the risks it poses
4. Develop effective governance, risk management and control processes for AI models to ensure they do what we want them to
5. Organize ourselves to properly steward the most important component of the current AI wave: the data
6. Strengthen our internal capacity in AI development
7. Ensure maximum public value for our investments when procuring AI technology from vendors
8. Put in place the right platform for development and deployment of AI solutions
9. Design a framework for monitoring performance and evaluating success of AI solutions to prove value to Canadians

1. Why an AI Strategy?

Artificial intelligence is receiving an unparalleled level of interest from just about everywhere (including governments, private enterprises, academia and citizens). At ESDC, this interest has manifested itself into several exciting initiatives that are expected to transform how we do business, paired with an equal amount of concern about the risks AI poses. There are questions that need answering related to what activities are underway, what plans are out there, and how we intend to ensure AI is incorporated into the department in a responsible manner that engenders public approval.

The AI Strategy, presented as an evolving, dedicated website, will outlay ESDC's plans, investments and current thinking on artificial intelligence. The Strategy will also provide the foundation to kickstart AI Governance and Policy, which will evolve as the department matures in this area. The Strategy provides an excellent opportunity for departmental staff to become up-to-date in our thinking around AI and enables a myriad of collaborative opportunities to push AI forward as a department to ensure we use it the right way.

Ignoring AI is not an option

One thing that is readily clear is that artificial intelligence is changing the world as you're reading this. From the explosion of chatbots that now represent the first point of contact for many client service journeys, to AI's that diagnose early onset of disease, more and more of our daily life is being augmented by machines. Public expectations are evolving right alongside.

Without substantive investments in artificial intelligence, ESDC's capacity to deliver services will lag further behind that of the private sector (and other governments for that matter), thereby eroding public trust. Service queue sizes, client wait times, client experience/satisfaction, outreach, policy analysis, research, internal services and management practices can all be improved through the use of artificial intelligence, and not embracing these potential benefits would be doing a grand disservice to the taxpayer.

1.2 What is Artificial Intelligence?

The Strategy Definition of Artificial Intelligence

For the purposes of this strategy, artificial intelligence will refer to digital solutions that exhibit human or higher-level judgement to carry out tasks for the department.

Further, artificial intelligence solutions must operate in one or more of the following active areas of AI development applicable to ESDC:

- Natural Language Processing
- Computer Vision

- Audio Processing
- Client Segmentation and Advertisement
- Strategic Optimization

Lastly, it is assumed that artificial intelligence solutions covered under the scope of this strategy contain at some machine learning elements that enable them to continually improve their ability to carry out the task.

Definition Implications

A primary objective of the AI definition is to identify which projects and initiatives will fall under the scope of AI Governance and Policy. The above definition is relatively tight compared to how the term is defined in other areas. Robotic Process Automation, for example, which often does not use machine learning in favour of simpler rule-based methods, is excluded. The objective behind this is to ensure early AI Governance activities focus on core artificial intelligence models that exhibit judgement/discretion in their predictions, as these models represent the greatest unknowns with respect to much of AI Policy. It will be, however, the role of AI Governance to institute a definition of AI that is appropriate with the finalized governance model, and adapt this definition as technology and the department evolve.

Other implications of any AI definition is that it will have inconsistencies with:

- Media (for which the term is often used to incite emotional response from viewership)
- Vendors (for whom there is great incentive to use the term for marketing purposes)
- Our clients (who have a range of views on the subject)

Much of the strategy is dedicated to empowering ESDC with knowledge to have informed conversations about many aspects of artificial intelligence, and many of these concepts apply to solutions that would fall under looser definitions of AI. Sound judgement is required from the whole department with respect to how far prudence outlaid by AI Policy will extend.

Further details about the current nature of AI, and their relevance to ESDC are presented in Section 3 of this strategy.

1.3 AI Strategy Objectives

What is the Strategy Aiming to Do?

The overarching objective of the AI Strategy is to launch official **department-wide** activities that pertain to artificial intelligence. As we begin, these activities will take the form of thorough conversations about the different facets of AI, and the Strategy aims to support this goal by providing sufficient information so that informed discussions can take place. The three main directions this Strategy outlays are the demystification of AI across

the department, the institution of AI Governance and the development of ESDC AI Policy, and positioning ourselves to deliver maximum value for the taxpayer in our AI initiatives.

Demystify Artificial Intelligence

A primary objective of this strategy is to provide an early resource for ESDC staff to gain greater insight into what modern AI is and how it works. Put bluntly, knowledge is power, and this has never been more true than during the information age. As awareness of current AI technology grows, more opportunities for incorporating AI components into ESDC business present themselves.

Further, a comprehensive departmental understanding of artificial intelligence will be critical to build, buy, support, integrate and leverage AI investments in such a way that the utmost care is taken with taxpayer dollars. To support this objective, a department-wide communications strategy for AI activities is being jointly developed by the Chief Data Office and the Innovation Lab.

Section 3 of the strategy presents the strategic direction for AI demystification. Therein, current plans for AI communication are presented. You will also find detailed discussions on different tasks AI can perform, and plain language descriptions of how they work. Also presented are some examples of how different types of machines can be applied in the ESDC context, with the objective of prompting new applications of AI within the department.

Set the stage for AI Governance and Policy

As we integrate AI into our work environments, it is essential that we consider the impacts (both positive and negative) and potential risks before unintended consequences worsen peoples' lives. As the department matures in its use of artificial intelligence, it will accordingly solidify its risk mitigation strategies through appropriate policy.

The development of AI Policy will be achieved by first putting in place a governance framework that includes expertise from the AI-pertinent functions of the department. Further, as many considerations of AI Policy are still in their infancy around the globe, AI Policy will be need to developed through moving AI projects through the governance framework such that specific policy decisions can be informed by real experience. Lastly, AI Governance and Policy will strategically embed itself into existing governing bodies and processes (e.g. Data and Privacy Committee, MPIB) where possible, such that new, standalone AI-based committees are kept to the minimum needed.

Section 4 details the strategic plan for an ESDC AI Governance by providing a proposed initial governance model, and some initial timelines for implementation. This section also presents some initial discussions pertaining to different aspects of AI Policy, highlighting the department's current state in these areas, along with anticipated challenges.

Position Ourselves to Deliver Value for Canadians

The ultimate goal of our artificial intelligence activities, as with any public service function, is to deliver maximum value for taxpayer investments. This will be achieved through

making the right internal investments to ensure we're not overly reliant on external providers, organizing ourselves to get better value out of our data, negotiating for vendor services in a way that provides flexibility in how the department can leverage them, and ensuring AI investments result in broadly used solutions that positively impact the lives of Canadians. Section 5 presents strategic initiatives that enshrine taxpayer value as the primary consideration.

Intended Audience

The audience for this document is departmental artificial intelligence stakeholders, including:

- Business and policy areas that seek to make to use of artificial intelligence
- Data science areas that enable AI activities
- Corporate risk management functions that are adapting to the new implications of AI

As mentioned previously, the strategy aims to provide a sufficient starting point such that well-informed conversations can take place. Accordingly, portions of this document will apply to different areas in the department to different degrees, yet the hope is that much of the material is written in an accessible enough manner for broader departmental consumption. Further communication efforts will of course better realize this objective into the future.

What the Strategy is not

The strategy does not go into any detail with respect to financial management or resource allocation, as AI can be incorporated into almost every function in the department.

The Strategy also does not detail specific technologies or tools that form part of the vision, as the rapid evolution of the AI landscape would almost certainly result in a futile exercise.

Grounded in Experience

Finally, the strategy isn't only the result of tireless, deep, policy thinking. The Data Science Division at the Chief Data Office has partnered with a number of different teams within ESDC to deliver on a range of different AI pilot projects. While these projects delivered value for the department in and of themselves, they also were designed to enable us to gain the necessary knowledge to ensure the AI Strategy is informed by proper experience.

Section 2 presents summaries of ESDC AI pilots, and anticipated future use cases that are expected to shape AI Policy as the department matures.

2. Grounded in Experience

The machine learning wave of artificial intelligence is relatively new. The level of understanding of recent technological developments vary from organization to organization. The AI vendor landscape is constantly evolving, with the industry not yet mature enough to comprehensively deliver commercial off-the-shelf products. There is no universally agreed upon ethical framework or rulebook for responsible AI development and deployment. In short, the AI issues we're trying to address have not yet been solved globally, and ESDC must put in its fair share of effort to support consensus on these issues.

Prior to and during development of the AI Strategy, it was recognized that our strategic direction needed to be informed by experience through AI pilot projects. These projects have delivered significant value to the organization in their own right, but have also enabled the Chief Data Office and other data science teams in the organization to:

- Dig deeper into the mathematics and algorithms such that inner workings of AI are both accessible and explainable to the department,
- Gain a comprehensive view of relevant existing policy, determine to what degree it applies to AI, and identify where AI-specific policy needs to fill in the gaps,
- Come face-to-face with the risks AI poses, and be able to articulate the challenges ahead,
- Inform the relevant areas of ESDC of what the department will need with respect to oversight, infrastructure and measurement frameworks, and
- Refine our view of what strategies we'll need to put in place and how we'll need to organize ourselves to deliver maximize public value for our AI investments.

The remainder of this section presents some examples of AI use in the department to date, and also presents a strategic roadmap for the types of use cases that the department expects to move into as it matures in its AI understanding.

2.2 Our Current and Future Priorities for AI Use Cases

With respect to the future, driving factors of how ESDC will mature with respect to artificial intelligence are the department's current plans related to service transformation and other major initiatives. There are a number of innovative activities underway already that plan to leverage AI, and it's important that efforts related to AI Governance and Policy align such that ESDC is ready to meet the demands when the time comes.

At the same time, as with any new area of investment, a great deal of prudence must be applied to ensure the department does not assume too much risk too soon. ESDC has thus far taken the approach of having its initial investments into AI focus on relatively safe areas that do not have unreasonable levels of downside risk for our clients or employees. However, there is long-term benefit to strategically pushing the boundaries of the

department's comfort level with respect to AI, as greater levels of service delivery can be achieved through this type of innovation.

The following presents a priority list of different applications of AI that represent a roadmap of how ESDC strategically plans to mature in AI, representing increasing levels of risk and reward as we become proficient, and AI Policy takes shape.

Areas of artificial intelligence that will be addressed by AI Policy in the short-term (2019-2020 - timelines not final):

- AI's that organize our work and triage workflow
- AI's that help our agents find information related to our programs and policies
- AI's that identify clients in our system with specific characteristics
- AI's that automate low risk business processes
- AI's that scan unstructured client data to populate structured databases
- Internal and external-facing chatbots

We know for certain that these activities will need to be addressed by early AI Policy as they are already being used in our day-to-day work environments.

Areas of artificial intelligence that will be addressed by AI Policy over the medium-term (Late 2019-2021 - timelines not final):

Additionally, other AI activities that are being actively investigated to assess potential are:

- AI's that monitor unstructured data for performance reporting
- AI's that inform our targeted outreach operations
- AI's that provide supporting information and recommendations to human agents rendering administrative decisions
- AI's that support policy analysis through emulation of the real world
- AI-based enhanced client authentication

Items planned for the longer-term:

Finally, more ambitious initiatives form part of the longer-term vision, but require a more mature AI environment at ESDC than what exists today. Examples include:

- AI's that automatically render administrative decisions in real-time
- AI's that draft documentation and/or communicate directly with clients about their file

3. Demystifying Artificial Intelligence

3.1 Communicating ESDC's AI Strategy

Public Servants are the primary audience for the ESDC AI Strategy. They can be broken in to two categories:

1. Public servants in oversight, decision-making, policy and enablement roles that need to understand the policy considerations of developing and deploying AI at ESDC.
2. Public Servants working in business areas that are looking to identify if AI can provide a good tool for dealing with their business problems.

This section of the strategy called 'demystifying AI' provides both audiences an initial understanding of how AI can be used to improve their business processes. The pilot projects highlighted in the draft further illustrate concrete examples of AI in action.

The CDO recognizes that further communications efforts are needed. We will work with the Public Affairs and Stakeholder Relations Branch (PASRB), and the Innovation Lab to create workshops, visuals and other documents so that all audiences can consume and reflect upon the strategy. The more people who are exposed to the strategy the more powerful and meaningful it will be, allowing it to be socialized both inside and outside the organization.

Once we are ready, an external communications plan will also be needed and should illustrate to the public that we are using AI responsibly through sound governance and policy. At this moment, we do not believe active communication to the public is necessary. However, we are always committed to transparency (the strategy has been ATIP'd) and will continue to answer media inquiries on the topic.

3.2 Background on the Artificial Intelligence Definition

Artificial intelligence as a term has no universally agreed upon definition. Its primary use in the most recent wave of excitement is in marketing and advertisement, with many developers having incentive to label their products as AI, even if their level of sophistication would have been state-of-the-art 30 or 40 years ago. Mathematicians, computer scientists, economists, science fiction writers and professionals from a variety of other fields all have different views on what constitutes AI, with some believing that the term should not be used at all due to its ambiguity. One of the benefits of continued use of the term for strategies such as this one, however, is in the level of interest that it generates, which in turn will help move forward the strategic initiatives outlaid in this document.

Artificial Intelligence has a temporal component

A notion commonly associated with artificial intelligence is that it is present in machines that replicate human judgement. This criterion alone as a definition of AI presents a scope issue, as computers can perform many tasks that are only otherwise doable by humans (for example, adding two large numbers together or identifying whether a word is present in a document). No one would consider these tasks artificial intelligence today, but at one point, they would have been considered the height of computing. It is therefore reasonable to introduce a time dimension into the definition of AI, and we will accordingly restrict our definition to refer to modern applications of machines that replicate human judgement.

Specifically, state-of-the-art research into artificial intelligence is predominantly concentrated in the following areas:

- Natural Language Processing: refers to the ability of machines to read, understand, categorize, summarize, extract information from and create information in written natural language.
- Computer Vision: refers to the ability of machines to classify, recognize patterns in and extract information from images.
- Audio Processing: refers to the ability of machines to listen to, infer sentiment from, extract information from and produce sound.
- Client Segmentation and Advertisement: refers to the ability of machines to analyze and predict patterns of human thinking and behavior, particularly in the area of online decision making.
- Strategic Optimization: refers to the ability of machines to, given a set of possible actions to take, along with a set of constraints defined by the environment in which they operate, make optimal decisions to achieve some objective.

All of these areas are relevant to ESDC, and specific details about each are presented in the remainder of this section.

What is machine learning?

Another term commonly linked with artificial intelligence is machine learning, which has a much more objective definition in the fields of computer science and mathematics. Machine learning refers to a set of statistical algorithms that enable machines to progressively improve performance at a specific task (i.e. to "learn") without being explicitly programmed on how to improve. Machine learning can be further broken down into three main sub-fields:

- Supervised learning, for which the algorithm is presented with available inputs and desired outputs, and programmed to learn itself the best relationship between the two so that it can predict outputs for future inputs.
- Unsupervised learning, for which the algorithm is programmed to find hidden structure in data, without explicitly being told what it is looking for.

- Reinforcement learning, for which the algorithm is not told explicitly what to accomplish, but rather given a reward signal based on its actions, and then has to interact with its environment to determine how to obtain the best rewards.

The current wave of artificial intelligence largely uses machine learning algorithms to develop solutions. Previously, the programming of solutions was much more explicit, leveraging experience of both programmers and relevant field experts. An enlightening example about the current wave of A.I. versus previous waves is perhaps in the analysis of medical scans to diagnose tumors or other diseases/conditions. Twenty years ago, an application developer would have sought the expertise of a top doctor in the field, and explicitly programmed the doctor's thought process in assessing the scan and diagnosing the patient, with some minor probabilistic modeling involved if the doctor assigned a level of uncertainty about any portion of his/her assessment. Today, the machine is given millions of images along with the patients' associated diagnoses by many human doctors, and programmed to determine the relationship (i.e. the "thought process") itself.

3.3a Text Classification

What is Text Classification?

Text classification is the activity of determining whether free text data (sentences, documents, etc.) meets certain criteria such that labels can be applied to it. Examples of this include whether a given text conveys a positive or negative message, if an email is a request for something or not, or if a given news article is potentially relevant to our department.

Text classification can be very useful to automate the triage and/or labeling of large collections of documents and where there is a constant flow of incoming data. In this context, the automation of the classification can not only free human resources for higher value work, but could also expedite the task such that new opportunities for action are presented.

How do Text Classification algorithms work?

Most text classification models today use supervised learning, where a statistical model is able to learn from human-labeled examples of how it should classify documents. For example, a text classifier learns what constitutes a positive message in a document by seeing documents that been previously labeled as positive by a person. The statistical model itself isn't told **how** to determine a positive message from other cases, but will instead use the example pairings of data and human labels to determine what it thinks is the best relationship. This model can then be used to predict/associate categories to future documents that it will have never seen before.

The mechanics behind text classification can be summarized by the following steps:

Estimation and Prediction Process

- To estimate a classifier, training data needs to be manually labeled with the desired categories.
- Feature development: To enable prediction based on the available text, the machine is given the opportunity to include a wide array of "features" of the text so that it can determine what is important to the classification. Features can include the presence of certain words, the presence of word pairings or triplets, presence of certain character patterns, part-of-speech tags associated with words, and many, many others. This step results in structured data (pertaining to the text) being created from the free text data such that it can be fed into mathematical models.
- In model training, the machine learns the optimal relationship between the available features and the desired categories. In modern text classification, this is achieved using deep learning techniques.
- Once these features are learned and an acceptable error level is reached, the model is ready to infer one or several labels/categories to text examples that it has never seen before.

How can Text Classification be used at ESDC?

Text classification can be applied very widely across ESDC due to the large amount of text data that is received and stored by the department. Examples include:

- Supporting back-office staff by triaging work items that arrive or are internally assigned in written format, enabling more efficient assignment/allocation of work and freeing agents for higher value tasks.
- Undertaking sentiment analysis to categorize the feedback of Service Canada clients as positive or negative.
- Triaging daily news articles, identifying when something is being said about the department, its programs or its Ministers. A text classification tool can be used to rate articles according to their level of importance or relevance to a wide range of professionals across the department.
- Helping with volume management for screening process to assess whether applications meet certain criteria. This could be applied to both internal assessment processes, and public-facing assessment processes.

What are some important risk considerations for Text Classification?

The main risk associated with text classification is that it will classify data incorrectly. Depending on the nature of the classification task, this could result in information not properly being transmitted, work being assigned incorrectly, or improper action being taken on client files. These risks would need to be measured against the error rate of an un-automated process, and the overall value the model provides withstanding these risks.

One useful feature when assessing classification risk is that models give a degree of confidence associated with their predictions, enabling a dimension of risk management (e.g. the automatic triage of the request will only occur if the machine is confident in its prediction above a certain threshold; otherwise the process will revert to the existing manual process). Further, as these models continue to learn from incoming data, AI models will improve their prediction accuracy, providing some assurance that the frequency of incorrect predictions will subside over time.

3.3b Information Retrieval

What is Information Retrieval?

Information Retrieval (IR) is the activity of obtaining the location of specific information from large knowledge repositories in response to a user need.

In the context of text information retrieval, modern IR mines information from the vast knowledge base and gives results based not just on keywords but also on the intent of the query. Additionally, it can take into account user personal preferences, different meanings of words and spelling errors. Information retrieval is also a critical building block for other AI applications, as it is commonly used in question answering / chatbots to retrieve passages that are likely to contain the answer to a user question.

A few popular applications using IR are the following:

- Search Engines (Google Search, Bing, etc.)
- Job Matching Websites (LinkedIn, Indeed, etc.)
- Google news

How do Information Retrieval algorithms work?

The mechanics behind information retrieval can be broken down generally in 3 major steps:

1. Preparing (indexing) the knowledge repository for more efficient information access.
2. Retrieving documents within the repository that match the information need.
3. Ranking the documents retrieved by relevancy to the information need.

For step 1, the purpose of storing an index is to optimize speed and performance in finding relevant documents given a query. Without an index, the search engine would need to scan every document in the corpus, which would require considerable time and computing power. For example, while an index of 10,000 documents can be queried within milliseconds, a sequential scan of every word in 10,000 large documents could take hours.

For step 2, the retrieval of the documents from the corpus is typically done to obtain a smaller set of documents that are relevant to the information need. This is typically useful to filter out irrelevant documents and to reduce the processing required for the ranking step.

For step 3, the purpose of the ranking of the documents is done to be able to assign an order of relevance to each document. The objective is to more highly rank documents that are more likely to contain the information the user is looking for.

How can Information Retrieval be used at ESDC?

IR at ESDC can be used in number of areas. The department holds many knowledge repositories of unstructured data from which it can be difficult to retrieve information (often the solution is navigable drop down menus, which require large amounts of both experience and patience). Also, since IR can be used as a component of other AI techniques, it can enhance other tools by providing relevant information in relation to user needs. The following use cases are possible with IR:

- It can be applied to HR needs in terms of retrieving candidates that match certain requirements.
- It can be applied to recommender systems in retrieving benefits that might interest specific clients based on unstructured data made available.
- It can be used as a component to enhance internal chatbot capabilities.

What are some important risk considerations for Information Retrieval?

Information retrieval generally carries less impact risk than other AI initiatives, as solutions usually provide suggestions to human users with respect to the location of the information they are looking for. Ideally, if the wrong information is ranked higher than the correct information, the user can manage the ranked results to find the particular clause they are looking for. However, ineffective IR tools are not riskless. In more serious cases, the IR tool may not present the correct passage as an option, and if the user is unaware, could proceed with downstream tasks based on incomplete information. IR tools can also be frustrating to use if they are consistently ineffective at providing the user the information they need.

This risk, of course, needs to be measure against potential risks of alternative solutions. Drop-down menu interfaces typically require significant investments in training to become an effective user, and more often than IR tools present the risk that the correct information will not be retrieved, especially for inexperienced users who do not know where to look.

Sources: - [Information Retrieval and Evaluation of the Privacy Risk on Twitter](#)

3.3c Question Answering (for Chat-bots)

What is Question Answering?

Question answering (QA) is a field of research in artificial intelligence that uses machine learning techniques to automatically answer questions posed by users. Questions are typically either typed into a chat window, or spoken aloud and then converted into text via speech recognition software. Responses provided by the machine extend beyond information retrieval (which provides a list of relevant passages from a knowledge

repository that a user can browse), and instead aim to retrieve the direct answer to the question by extracting it from the most pertinent location. The answers are then provided back to the user by presenting response text on the screen, or converted to speech using text-to-speech software.

Examples:

Q: What is the maximum monthly OAS pension one can receive in September 2017? A: \$583.74

Q: Who is eligible to sign an EI Sickness medical certificate? A: A medical doctor or other medical practitioner (health practitioner)

How do question answering algorithms work?

The mechanics behind classical QA algorithms differ across organizations, but generally follow these major steps: (reference Jurafsky & Martin)

1. Based on the question, determine what type of response the user is looking for (e.g. location, regulation, and date).
2. Generate a query based on the wording of the question and the response type identified in (1).
3. Rank documents and databases in a knowledge repository by relevance to the query generated in (2) using information retrieval techniques, then rank passages in the highest ranked documents again by relevance.
4. Use information extraction techniques to generate a list of possible answers, rank them, and provide the highest ranked response(s) back to the user.

The knowledge repositories used in these algorithms differ based on function. General purpose QA tools use the Internet as the main repository, whereas organization-specific QA tools will use organization-specific documentation sets, which are often proprietary and not available online. Many of the sub-algorithms developed in the major steps listed above are not as general as those for information retrieval, and therefore more work is required to adapt them for other purposes. However, recent progress has been made in this area using modern deep learning techniques, which is enabling models to become more and more effective at successfully answering questions using broad, general-purpose QA databases (reference Squad).

How can question answering be used at ESDC?

There are several potential applications of question answering tools that could readily provide value to the department, serving both Service Canada clients and ESDC portfolio staff. Examples include:

- A QA tool available to the public that answers general questions related to ESDC program eligibility, application procedures, available benefits and other important

information for which clients would otherwise have to navigate through sub-menus of websites to find the information themselves.

- A QA tool available to support front-line staff during client interactions. The tool would again prevent the need for staff to navigate through list-based knowledge repositories to find the information they need to serve the client, thereby reducing service times.
- A QA tool available to all ESDC staff on the Department's intranet site to support personnel in their day-to-day work. The AI could answer a broad range of questions from "What software is departmentally available for process mapping?" to "How do the CPP child rearing and credit splitting provisions interact?" to "Who is responsible for managing the Job Match algorithm at ESDC?"
- A QA tool to support policy analysts, researchers and program officers for which the knowledge repository is based on reports, data tables and other sources that form the evidence base for their work. Such a tool, for example, would answer questions like "What is the current youth unemployment rate in Chatham?"

The first example, and to some degree the fourth example, utilize public facing information, and could consequently be developed independent of ESDC investment. Google Assistant, as one example, is becoming increasingly effective at general QA over all internet sites, and will naturally increase its capability to answer ESDC related questions as time progresses. The latter two examples include knowledge repositories that are restricted to departmental employees, and would therefore need to be developed internally or through an external contract that provided the repository to the vendor. Additionally, though separate tools are listed above, ESDC could also implement a single multi-purpose QA tool that restricts or expands potential answer types and knowledge repositories based on user type.

What are some important risk considerations for question answering?

The main risk inherent to QA tools is that they will inevitably answer questions incorrectly. This in turn could lead to misinformed clients or staff, who then take inadvisable decisions based on that information. However, competing sources for information retrieval also generate incorrect answers (clients misinterpret information on the web page, or a misinformed staff member passes on incorrect information to another now-misinformed staff member). As time progresses and algorithms continue to improve, it is also inevitable that the likelihood the AI provides the correct information will surpass that of other methods (if not true already).

From a personal information perspective, so long as the knowledge repository does not contain any personal information, there is no privacy risk associated with the associated QA tools. Controls of who can access what components of specific knowledge repositories should be managed independently of question answering front-end applications.

3.3d Text Generation / Document Creation

What is Text Generation?

Text Generation refers to algorithms in the domain of AI and machine learning that can produce writing in natural language (i.e. they generate their own text). Text Generation is currently in its infancy, but there are already a variety of different real-world applications which provide a glimpse of its potential. Generative algorithms can be used to create texts of arbitrary length, such as a short poem, a description of an image, computer code, or even a full-length novel.

Example: A machine learning algorithm can learn Shakespeare's writing style and begin to generate text that to most observers mimics the language used by Shakespeare. Some compelling examples of the potential of text generation can be found [here](#).

Text generation algorithms are part of a broader family of generative algorithms that are used to generate images, audio, and other media. Many neat examples exist with people [having fun with generative models](#).

How does Text Generation work?

There are several different frameworks for generating text (and generative algorithms more broadly). One class involves training algorithms to predict the next word in a sentence, or character in a word. With this type of model, one feeds in input texts, and the algorithm learns to predict the next word or character in a sequence of words or characters based on word/character sequences the machine has observed in large collections of text. Once the model is trained, one can generate text by feeding in input and then having the algorithm predict the next item in the sequence while recursively feeding in the algorithm's output as new input.

Up until recently, a class of deep learning models known as Recurrent Neural Networks (RNNs) were the most effective algorithms at generating text in this fashion. RNNs, and in particular a variant known as Long Short-Term Memory Networks (LSTMs), have the ability to hold previous inputs in memory, making them ideal for a problem in which the task is to predict the next item in a sequence. A drawback of this type of method is that the machine doesn't incorporate any semantic meaning or context associated with the text; it is simply repeating patterns it has previously observed.

An entirely different class of generative algorithms are known as Generative Adversarial Networks (GANs). GANs work by pitting two distinct models against each other, one called the generator, and the other the discriminator, in an adversarial fashion. The basic idea is that the generator network will attempt to generate data that mimics real data (in this case text), while the discriminator network is trained to determine which data is authentic and which data was generated by the generator network. In this way the generator network must learn to fool an increasingly effective discriminator network. For example, one could train a GAN to generate news articles by having the generator network attempt to generate news stories while the discriminator network attempts to distinguish between real news

stories, and the ones that have been generated by the generator network. Both the generator and discriminator network must be machine learning models themselves, so it is natural to use an RNN based architecture for the generator network for the same reasons outlined above.

Most recently, a class of text generation models that have shown an uncanny ability to produce original text use a feature called "attention" (reference paper). Attention models enable machines to maintain context for lengthy passages within text, consequently providing machines the ability to write/speak for longer stretches about a subject. As this is a very active area of research, breakthroughs are happening at a rapid pace, and ESDC must exert the effort required to stay up-to-date, both for the exciting opportunities generative models present, and for the alarming risks they pose.

How can text generation algorithms be used at ESDC?

Text generation has some near-term potential for application at ESDC, as well as significant potential for applications in the medium to longer-term as generative algorithms improve.

In the nearer term, text generation can be used to build more sophisticated chat-bots and to generate summaries of documents. Sophisticated chat-bots, such as the new Google assistant that has emerged from the Google Duplex project, utilize speech-to-text, text-to-speech, and generative models to generate both text, in order to formulate responses, and also audio, in order to generate a voice to converse with the individual who is engaging with the chat-bot. This kind of chat-bot could be used to improve Canadians' experience in obtaining information regarding ESDC's programs and services, as well as providing a way to automate reaching out to Canadians if ESDC needs information, or to notify them that they are eligible to apply for a program and provide information on how to complete the application process.

In the longer-term, assuming that there are significant advances in generative algorithms for text generation, this technique could be used to generate drafts of briefing notes, summaries of documents, and even presentations. There has also been some limited success in having algorithms generate code in multiple programming languages. Generative algorithms have thus far been more successful in the realm of audio and images than text, but it is more difficult to envisage the applicability of these techniques to the department, excepting the aforementioned voice generation for chat-bots.

Another example of how text generation could be useful is to help an algorithm explain the decisions it makes. It is often challenging to interpret and understand why a machine learning algorithm performs the way it does, but adding a generative component to algorithms can allow them to generate text explaining why the model output the way it did, in some ways analogous to human beings explaining their reasoning. This is currently not possible beyond some very limited and rudimentary examples, but there is active research being undertaken.

What are some important risk considerations for generative algorithms?

Generative algorithms pose a serious risk for society as a whole, as they enable individuals to create fake audio, images, video, as well as sophisticated bots online, and, eventually, likely convincing fake news articles. There is also a risk that generative algorithms will start to replicate undesirable behaviour. As with all machine learning algorithms, generative algorithms learn from the data they are trained on, and if poor judgement is reflected in the training data, the algorithm will learn to replicate it. For example, Microsoft released a chat-bot that was trained via interaction with the public, but after conversing with different individuals, the bot started replicating hate speech.

Generative algorithms would also pose some risk internally here at ESDC. These risks include a chat-bot providing incorrect information to Canadians, a generative algorithm producing incoherent content when communicating either internally or with a member of the public, or the reproduction of undesirable behavior as mentioned above. These risks can be mitigated by extensive testing and by limiting the scope in which the algorithms are used to areas in which mistakes will not be overly consequential. Careful controls on the training data used by the generating algorithm are also important to ensure what the algorithm generates remains acceptable.

It is also important to note that if generative algorithms allow us to provide a service that we are not currently able to provide to Canadians, it may not be essential that they achieve a high degree of perfection - as long as they are improving the services we are able to offer, we can tolerate a certain degree of error or imperfection.

3.3e Computer Vision

What is Computer Vision?

Computer vision refers to algorithms that gain a sophisticated understanding of the content of digital images or videos (sequences of images), enabling tasks such as image classification, object detection, scene reconstruction, video tracking, and many others. Input images can take different forms, such as standard 2-dimensional images, image sequences, and 3D medical scan images. The essential element is that digital images have a finite set of digital values (associated with "pixels").

Enormous progress in computer vision has been made worldwide in recent years. Applications include photo-tagging, self-driving cars (for which cameras are constantly recording the vehicle's surroundings to enable it to make decisions), and even art generation. Specific individual applications are very narrow in scope (meaning each model is developed to perform one task in particular), while a general computer vision AI that can completely understand and interpret the full contents of an image to the same degree humans can is beyond current technology. However, at individually targeted tasks, computer vision AI has surpassed human capabilities in a number of areas.

Note that Computer Vision differs from image processing, which is generally associated with image editing and formatting tasks such as image restoration, digital enhancement

and segmentation. Image processing usually results in the manipulation of one image to yield another, while computer vision is more interested in creating structured data from an image (e.g. Is this feature present? Where is it? How big is it?) that can be actioned for other tasks. Both classes of techniques rely heavily on machine learning in modern applications.

How do Computer Vision algorithm work?

Most computer vision applications are developed in a supervised learning context (i.e. they learn from examples). One classical example is a model that is able to recognize if an image contains either a dog or a cat (an example of image classification). To achieve this, the model will be fed millions of pictures of dogs and millions of pictures of cats (labeled accordingly), and teach itself to learn what specific pixel combinations are useful for predicting whether a dog or cat is present. Once trained, the model will be able to make incredibly accurate predictions on new images that it hasn't seen before.

Object detection algorithms, a step beyond image classification, identify where, in an image, a particular object sits. Similar to image classification, most object detection algorithms also use supervised learning, where the model is provided with millions of images along with pixel locations of various objects, and the machine trains itself to localize objects of interest.

Computer vision algorithms typically are provided an image as input, and asked to interpret that image. A traditional pipeline CV applications might look like:

- Image Processing / Manipulation
 - Image Editing
 - Image Enhancement
- Computer Vision Algorithms
 - Image Classification
 - Object Detection

How can Computer Vision be used at ESDC?

As the department still uses physical paper for many processes, there is significant opportunity to use computer vision to render the physical files into machine readable formats and structured database. These applications would primary use an area of computer vision known as Optical Character Recognition (OCR), which aims to retrieve text from images and videos.

Example applications of computer vision in this context are:

- Computer vision can be used to enable the automatic extraction of typed text or hand-written text from documents submitted by clients (e.g. forms, passports, ID cards).
- Computer vision can be used to detect potential fraud attempts in physical documents that the department is receiving when clients apply for benefits.

Examples without text include the use of computer vision to better estimate wait times and queue sizes in Service Canada Centres with video tracking, or for client authentication via facial.

3.3f Client Segmentation

What is Client Segmentation?

Client segmentation (or market/customer segmentation) refers to dividing clients into groups based on common characteristics so that effective and directed client integration can be taken. For example, resulting segments might identify clients with high needs versus those with low needs, engaged versus disengaged, informed versus uninformed and so forth. In the artificial intelligence context, this activity relies heavily on different types of data/information (e.g. demographics, behaviours, file activity).

Segments can be used for targeted interventions, such as special outreach, directed marketing or policy making, with the aim of intervening before problems arise or identifying groups that are not adequately reached by communications or services. Modern segmentation takes advantage of groups directly identified from the data, which may show unexpected relations between certain groups of people. This is in contrast to 'bought' groups that predetermine which types of people are expected to behave similarly or have similar needs (e.g. youth under 25, urban women with children age 29-45), and may miss more nuanced connections.

In modern AI applications, client segmentation is taken to personalized levels (also known as "segment-of-one"). Using vast amounts of data and deep learning architectures to enable complex structures associated with client behaviours, their likes and their dislikes, AI's are becoming increasingly effective at gaining client attention and prompting response.

How does Client Segmentation work?

Data mining techniques identify groups of clients based on factors, traits and behaviours that are measurable in available data. Computational methods are replacing the older "business rules" approach that relies on perceptions of business experts, which can be biased/narrow in focus and unscalable. The objective of AI client segmentation in the modern context is to identify patterns of behaviour at such a fine level, that predictions about behaviour become different from person to person.

A central theme to modern client segmentation methods is that the data does the talking, almost always using a form of unsupervised learning. Every piece of data available on clients is fed into one or more unsupervised learning algorithms, and the machine is instructed to learn "hidden" patterns:

- For which characteristics or behaviours are groups of clients comparable?
- For which do they differ?
- What types of behaviours correspond to other types of behaviours?

Related techniques such as social network analysis are additionally brought in to help paint the client picture through their known relations. Natural language processing is leveraged for data on the client that is stored in free text format.

Naturally, as the desire of AI segmentation is predict behaviour at levels as close to individual as possible, data pertaining to a massive number of clients are required for the solution to become effective.

How can Client Segmentation be used at ESDC?

ESDC is responsible for a range of services affecting a very large and diverse set of clients across Canada, and segmentation can aid in improved delivery of those services and understanding of our clients:

Efficiency: Proactively find unreached clients for communications/services/programs

Accountability: Evaluate whether particular needs are met in specific client bases

Legitimacy: Understand client groups and their needs in policy-making and program delivery

For example, our *Poverty Reduction Strategy* considers different measures of poverty, since it is not homogenous and relates to different segmentations of people depending on how it is measured. Advanced analytics and segmentation methods can test assumptions about groups, for instance, to find which similarities matter for poverty reduction or which people are out of communication or service coverage with our current policies. Segmentation can help target needs not adequately met because of lack of understanding of characteristics and behaviours, and hopefully result in better service to Canadians in need.

What are some important risk considerations for Client Segmentation?

Certain challenges arise when segmenting individuals for targeted responses, these include:

- The quality of data will be directly reflected in the quality of a segmentation process, mirroring the distribution of individuals in the data. Questions of bias and under-representation in the data must be addressed at the outset.
- Collecting and accessing demographic/location-based/behavioural data (e.g. administrative data, market/satisfaction surveys) is time-consuming and expensive if not already collected for other purposes. Further, the use of all of this data for this purpose may present privacy and/or consent issues.
- Segmentation is inherently discriminatory, since it divides people for specific messaging/engagement. This activity can risk ESDC's reputation if we are not careful to correct bias (e.g. reaching the under-reached) instead of creating further bias. Further, unless properly monitored and appropriate counter-measures put in place, machines may start to segment on traits or behaviours in such a way that it violates the Charter of Rights and Freedoms, Employment Equity legislation or other human

rights protections (for example, segmentation based on skin colour or medical behaviours).

- Finally, many segmentation algorithms are not very interpretable, which often can render the machine's decision difficult to explain.

3.3g Strategic Optimization

What is Strategic Optimization?

Strategic optimization is an area of artificial intelligence within which machines seek to optimize their actions to most effectively achieve an objective. The most reported on types of strategic optimization AI's developed recently have been machines that play games and machines that trade stocks. In the former case, the most famous game playing machine is Alpha Go (which plays the strategy game Go), developed by Google's DeepMind. Here, the objective is to win the game subject to the game's rules, and the machine needs to determine its optimal move set to achieve that objective the most often possible. In the case of trading stocks, the objective would (obviously) be for the machine to make as much money as possible given a specific risk tolerance.

Unlike most applications of natural language processing, computer vision and audio processing, strategic optimization AI's are not necessarily restricted to unstructured data (though they do often use it). Another notable feature that separates strategic optimization from other areas of artificial intelligence is that, in many cases, machines aren't simply capable of replicating human judgement (or the best of human judgement). They are capable of surpassing it, in some cases to a degree which if a human trained his/her entire life, they still could not compete. Alpha Go has no issues dominating top professional human Go players, and high frequency stock trading algorithms have replaced day-traders across the globe.

How do strategic optimization algorithms work?

The most successful strategic optimization solutions are primarily developed using reinforcement learning algorithms, though are supplemented by supervised learning to help the machine learn more generally.

Reinforcement learning algorithms can be thought of as repetitive trial and error. The machine will start out by taking effectively random actions (given the situation it finds itself in), and through the reward signal it receives associated with those actions, begin to value taking certain actions over others. After repeated and prolonged exploration of which actions yield the best rewards, the machine will eventually discover which actions are optimal such that it can maximize its objective. Additionally, despite learning continuously that certain actions are superior to others, it will still explore what it believes are sub-optimal actions to reinforce whether or not the action is indeed sub-optimal.

In all real-world reinforcement learning problems, the ability of the machine to explore every possible situation it could find itself in (and learn the optimal action associated with

that situation) is impossible, as there can be an infinite number of such situations. In modern applications, supervised learning algorithms are used to help the machine generalize situations to those it has seen, in order to make hopefully optimal decisions for new situations it encounters.

Often, business problems in which strategic optimization is being applied does not lend itself well to real-world trial and error (e.g. situations where someone's health or financial livelihood may be at stake). In these situations, a simulated environment that is representative of the real-world is created such that the machine can experiment and learn. Once it becomes sufficiently proficient, it is then empowered to take strategic decisions in the real world.

How can Strategic Optimization be used at ESDC?

A field of study that lends itself well to strategic optimization AI is that of Operations Research, where the goal is generally to minimize time (or another resource) dedicated to a set of tasks. Service Canada is one of the largest operational environments in the country, and decisions are made every minute of every day at in-person centres, call-centres and back-office processing centres on how best to allocate resources to minimize client service times. There are numerous opportunities to explore the use of strategic optimization AI's to improve the efficiency of our operations.

Other longer-term potential applications of strategic optimization AI's could include:

- Virtual assistant bots that learn to automate our menial tasks so that we can focus on higher value work.
- The allocation and assignment of computing resources that are shared among users across the department, to minimize both machine downtime and user wait times.
- Malware protection and other IT security functions.
- Program outreach, where the machine learns what is and is not effective in promoting our programs to potential clients.

What are some important risk considerations for Strategic Optimization?

If the machine is given full power to "explore" its options, there is a potential danger in strict trial and error deployments that they model may explore a catastrophic action (e.g. stop payments to all Canadians). In many cases, it is important the machine be given a safe environment to explore its action space, and that the machine's action possibilities and constraints be carefully designed to ensure it operates within desired boundaries.

Relatedly, designing simulation environments can be very resource intensive, and thus cost-benefit analyses are prudent before investing heavily in strategic optimization learning.

4. AI Governance & Policy at ESDC

So far, the AI Strategy has articulated the great opportunities that AI presents for ESDC to transform the way we deliver services to Canadians. However, it also presents some very significant risks that need to be managed. Some of these risks are legal in nature; it will be important for ESDC to be clear on what it is allowed to do in the area of AI within the current legislation, and how it needs to protect itself legally. Others are more ethical; the department will need to have a sound understanding of the degree to which Canadians are comfortable with the use of AI in government and how bias will be controlled. Still more are logistical: even with sound AI strategic thinking, how do we ensure our AI solutions perform the way we expect them to, and we maintain control. All of these issues are compounded by the reality that public perception, technological advancements and other facets of AI are evolving extremely rapidly, meaning that answers to these risk considerations are often moving targets. These issues have not yet been resolved globally, so answers will not come easy. However, there are many organizations developing similar policy in this area, so we are not alone.

This section of the AI Strategy outlines how we intend to address these matters through the implementation of *AI Governance* and the development of *ESDC's Artificial Intelligence Policy*. ESDC's AI Policy will be developed through AI governance and will guide decisions in a number of areas related to acceptable use cases, transparency, accountability, privacy, security and model quality, among others, with its ultimate objective being to provide the public with confidence that we are going to use AI responsibly. We are not sure where we will land, so we will need to use pilot projects and a governance framework to figure it out.

The development of ESDC's AI policy will take place over the 2019-2020 Fiscal Year, but several initiatives are helping to lay the groundwork:

- The Treasury Board Secretariat [Directive on Automated Decision Making](#) defines requirements that departments must ensure are met in order to use Automated Decision Support Systems that provide external services. Though systems that provide external services only represent a portion of areas of applicability of AI for ESDC, they do represent the most critical area for AI Policy to protect Canadians. In this respect, the Directive represents a minimum level of care that needs to be applied to use AI. However, for ESDC we'll need more concrete policy to support our AI initiatives and AI Governance framework that extends beyond external services to include internal services.
- The ESDC Data Strategy has initiatives underway focusing on data privacy and security, and many aspects of privacy associated with AI are more related to who/what has access to underlying data (as opposed to what is being done with it). The ESDC Data Policy is currently being developed and will continue to inform the AI strategy.
- AI governance is taking shape with increasingly ambitious pilot projects to establish answers to AI Policy questions. As AI projects move through the governance process we will continue to note key AI policy dimensions that require specific focus. It is

important to keep in mind that the overall goal of the deployment of AI solutions is to improve service for Canadians. AI policies that result in a development process being more cumbersome defeat the purpose. AI policy consequently needs to be grounded in pragmatism, and final policy decisions will be properly field tested with appropriate test projects.

AI Policy will be a critical piece of the AI puzzle, and it's important that it be a joint effort by the department to ensure it meets the needs of all stakeholders. Consultations across the department in the 2018-2019 Fiscal Year have provided useful insights to this updated version of the strategy, many of which are discussed in the AI governance and dimensions of policy considerations.

The CDO will continue to engage across the department to implement AI governance and have an initial AI policy in place by the end of Fiscal Year 2019-2020.

List of Relevant Policy Links

- [Department of Employment and Social Development Act](#)
- [Treasury Board Secretariat Directive on Automated Decision Making](#)
- [ESDC Data Strategy](#)
- [Draft ESDC Data Policy](#)
- [Old Age Security Act](#)
- [Employment Insurance Act](#)
- [Canada Pension Plan](#)
- [Privacy Act](#)
- [Open as a Foundation for Digital Government](#)
- [Gender-Based Analysis Plus](#)
- [Financial Administration Act](#)
- [EU Guidelines for Trustworthy Artificial Intelligence](#)

4.2 Proposed AI Governance

Governance of Artificial intelligence at ESDC is needed in order to establish processes of decision-making among stakeholders involved in developing and implementing AI solutions. The objective of AI governance is to create a system of trust so the department can move forward with developing and procuring valuable AI solutions confidently, knowing that considerations related to privacy, security, ethics, law, transparency, bias and performance are made along the way.

This strategy presents a draft preliminary governance design that will be implemented to guide AI projects. Keeping with the spirit of being grounded in experience, policy will be developed by projects going through the governance process, where consensus will be built over time by decision makers interacting collectively to solve problems. This approach is necessary because AI is a platform technology with many possible applications and various risk profiles; it should be governed with an incremental risk-management approach that is case and context sensitive that is refined as we learn.

The design of the proposed AI governance process below aims to:

- Enable research and experimentation at the outset with low governance barriers, promoting an environment where it is ok to 'fail fast'
- Ensure comprehensive risk mitigation strategies are in place when we move from experimentation to development ("trying" to "doing")
- Leverage and embeds itself within existing governance structures where possible
- Include a lens that emphasizes value to clients, the organization and taxpayers (it's not just about risk aversion)

The design of the governance framework is formatted as a checklist of questions all AI projects will need to answer in order to provide assurance that the solution has been created responsibly. There are three main phases: a research and exploration phase; a design and development phase and an implementation phase. The end of each phase includes a milestone point where review and final decision making takes place by relevant stakeholders. This governance framework is designed to be agnostic with respect to internally developed versus procured solutions, providing flexibility to business areas in how they wish to develop their AI's.

Research and Exploration

The governance process begins with the objective of enabling the department to openly explore the possible. Business areas are encouraged to innovate and investigate how AI can be applied to a particular set of data to determine if it can improve a business process or solve a problem.

A small check with AI Governance is proposed to:

- Ensure the department has client consent to use their data for this purpose

- Make available tools and algorithms to support the exploration, such that existing work can be leveraged
- Confirm the project is aligned with pre-defined high-level AI ethical principles

After that, it will be up to the respective business lines to determine the feasibility of the solution and whether there is sufficient business value to proceed to solution development.

Phase	Objective	Sample Questions	Who sets the tone?
Initial Governance Check	Quick check to get started	Do we have consent to use this training data for this purpose? Is this type of AI model even feasible given current technology? Has it already been done? Project in line with high-level ethical principles & values?	DPC: Privacy Management Chief Data Office AI Ethics
Research & Exploration	Determine if solution can be built	How might the AI fit into the business process? Is the data of sufficient quality to use for this purpose? Does the data provide the predictability we want?	Business Area (with support from Legal, privacy, CDO, IITB as needed)
Business Review	Should we build solution?	Will the AI provide value to the organization?	Business Area

Design and Development

During the development stage of an AI project, relevant stakeholders (legal, privacy, CDO, IITB, and CFOB) will be involved so that the business area can keep the AI Policy dimensions of considerations top of mind as they build. Considerations will include model performance, controlling for bias, IT security, privacy, program integrity, accessibility, ethics and compliance with legislation, policy and directives such as the TBS Algorithmic Impact Assessment. This approach will foster a culture of enablement rather than prevention by allowing the business area to bring decision-making functions into the process early on rather than waiting until the review stage to find out that they should have considered something critical. In addition, this approach allows decision makers to make informed decisions about project risks and if the AI solution should proceed to implementation. To ensure we are pushing boundaries an acceptable and healthy failure rate should be established.

Phase	Objective	Sample Questions	Who sets the tone?
AI Design & Development	Design & develop AI solution with stakeholders	How do we measure and control for bias? How do we design the solution while safeguarding assets and information? How do we design the solution to ensure it is compliant with relevant legislation and policy? How do we consider and measure the potential impacts on clients and the dept.?	Business Area (with support from Legal, privacy, CDO, IITB)
	Answer question "Should we use AI for this?"	How will decisions made by the solution be explained? How will the solution uphold program integrity? How will the solution be compliant with the AIA and other TBS standards?	Above AI Stakeholders become involved to help ensure project success
AI Governance Review	Should we implement?	What legal risks are we assuming when using AI for this purpose, and how do we mitigate them? Are there any ethical risks to consider and mitigate? Is model quality acceptable for the business process? Are we comfortable with anticipated model bias and explainability? Does the AI's design meet departmental security standards?	DPC

This phase of the review framework will be embedded in existing governance processes (e.g., MPiB, EARB)

Implementation and Monitoring

If the decision is made to move the AI solution to implementation, relevant stakeholders will be tasked with making it happen. The objective of this phase is not to determine if the solution will be implemented, but rather how will it be delivered and how will we measure its success over time. With all the considerations made in the development process, the goal is to have close to 0% of projects fail in the implementation phase. The business area will work with the stakeholders to complete all the assurance initiatives needed to control risks and monitor success over-time. IT will play the main role in enabling the solutions deployment by setting up infrastructure, accessibility and data management support. In addition, AI specific issues and monitoring support will be provided as needed by other areas to track changes to the model over-time and protect against reverse engineering.

Phase	Objective	Sample Questions	Who sets the tone?
AI Implementation	How do we deliver the solution?	Will the code for the solution be made open and how? How will data from the solution be collected and stewarded? What infrastructure will support the solution? How will the solution be protected from reverse-engineering? How will the model be peer-reviewed and evaluated? How will the solutions performance be monitored? How will the solution meet accessibility standards?	IITB Business Area Other AI enablers
Review & Monitoring		Is there a monitoring and review cycle in place for the solution? Is there a recourse plan if the solution needs to be changed or removed? Is there a process for reflecting policy or legislative changes?	IITB Business Area Other AI enablers

Preparing our existing committees for Artificial Intelligence initiatives

A separate committee solely dedicated to supporting AI Policy decisions is not a preferred option, given the existing governance structure in the Department. It is felt that existing governance bodies and structure will be appropriate for AI initiatives and have been highlighted in the governance design:

- The Corporate Management Committee (CMC) is expected to take executive decision on all high-level, contentious AI issues.
- The Data and Privacy Committee (DPC) is expected to take executive decision on most policy aspects related to AI. The DPC, for the foreseeable future, will act as the default body for AI and Data Strategy issues that don't naturally fit with other governance committees.
- The Enterprise Architecture Review Board (EARB) will govern the IT aspects of AI, including architecture and infrastructure governance, setting technology standards and providing IT investment direction.
- Major Projects and investment Board (MPIB) will keep with its mandate of supporting rigorous and transparent project planning, project management, and investment decisions by allowing the AI governance stakeholders to provide assurance in the decision making process.

As an interim measure while AI Policy and Governance are being developed, the Data Science Division at the Chief Data Office will continue offer its services to groups within the department seeking compliance support for its AI initiatives.

4.3 AI Policy: Key Dimensions of Consideration


This section presents different ethical dimensions that are expected to be addressed by AI Policy. This does not necessarily represent an exhaustive list, but it is grounded in the experience gained from the AI pilots undertaken to date, consultations and feedback from across the department, and information obtained from thought leaders through research, conferences and forums on artificial intelligence.

We as a department will develop processes and standards for AI as its capabilities advance and change. The development of a comprehensive AI Policy relies on standards created by TBS, across the Government of Canada, in review of policies implemented by other countries (such as the steps taken in the EU to protect and educate citizens about AI and the use of data).

There are many facets to the ethical questions surrounding machine learning and AI. However, AI solutions may often not be worse than current operating procedures in terms of bias, privacy, or security, and have the potential to improve on these aspects. Responsible AI use can create value in tasks that cannot be replicated by hand. AI Policy cannot lose sight of this.

This section of the strategy discusses key components of AI policy; considers the risks and opportunities, and identifies next steps in their development.

4.3a Legal



Current Requirements and Policy

All departmental activities (AI included) need to comply with the [Department of Employment and Social Development Act](#), the [Old Age Security Act](#), the [Employment Insurance Act](#), the [Canada Pension Plan](#), and other relevant legislation.

TBS Directives such as the [Directive on Automated Decision Making](#) also prescribe actions with which the department will need to be compliant. TBS policy represents a minimum level of care that needs to be applied in the use of AI, and provides a reference point for departmental AI initiatives.

Considerations and Risks

ESDC AI Strategy

DRAFT - V3.0

Three main categories of legal consideration have been identified when the department decides to use AI:

1. **Choice of solution:** It is important that the choice to use an AI solution is appropriate for each associated business problem. For example, if the training of the AI system requires large amounts of data, but too little data is available, or not enough Canadian data that may reflect our laws and unique processes, then the AI solution may not be viable. Legal risks can compound if an inappropriate solution is chosen in the first place, and a "wrong" solution used that results in adverse client decisions is the fault of the department. Put bluntly, "AI for the sake of AI" presents significant legal risk.
2. **Build of the solution:** The accuracy and performance of an AI solution must be appropriate for the task it is assigned to do, so that vulnerable people do not get hurt. Model accuracy levels must be acceptably met prior to model productionization, taking into account the full business context and existing processes. Further, if an AI is feeding into a decision, then the Crown is liable for that decision (the software cannot defend itself). It is critical, therefore, that *ESDC maintain an adequate internal expertise that understands and is able to explain the inner workings of its AI models*. This also has important implications for procured vendor solutions, since vendors do not take on the responsibility of behalf of the department. This has the implication of outright excluding AI solutions that are "black box" for which ESDC staff have limited access and understanding for many business problems.
3. **Delivery of the solution:** Even if a solution is built as accurate as possible for an appropriate business problem, it still must be delivered, maintained and have an appropriate plan for foreseeable roadblocks (such as system downtime). The responsibility requires that we have solutions that can provide correct information but also communicate that information properly. If components of a system are out of commission, can we still explain an individual decision? Has sufficient risk mitigation been properly built in to the AI pipeline from early on in the design stage to post-production system measurement?

Current Status and Plan

ESDC's Legal Services Branch (LSB) has been very active in the area of artificial intelligence over the past several months:

- An [RFI](#) was launched in 2018 to procure private sector expert feedback on what AI can do, key considerations when implementing solutions, and other key considerations.

-

A pilot that warrants specific mention is a collaborative project between the Human Resources Branch, LSB and the CDO that is exploring the use of artificial intelligence in supporting the initial screening of candidates for a competitive employment process. This project, relative to our other early AI projects, has the potential to impact people's lives to a

greater degree, and was hence chosen as an excellent candidate to explore boundaries and build/refine appropriate AI Policy during project development.

LSB is expected to play a critical role in the development and implementation phases of the AI Governance process. A legal lens is essential for robust AI Policy that protects the interests of the department and its clients.

4.3b Privacy

The power of current AI technology is in the data, and almost universally, the more data available the better the AI will perform. However, many concerns exist relating to individual privacy in the collection, storage and use of personal data. AI Policy will need to be reflective of both the relevant legal privacy requirements and government policy objectives, as well as public opinion on appropriate data use, as key requirements for public trust.

Current Requirements and Policy

From a privacy perspective, AI Policy will be compliant with the Charter of Rights and Freedoms, the Privacy Act, the Department of Employment and Social Development Act (DESDA), Treasury Board policies and directives and ESDC's policies, specifically, the Departmental Policy on Privacy Management (DPPM). These policies, legislation and directives are the most relevant law pertaining to appropriate collection and use of data for the department. The notion of consent (users being appropriately notified and provided the opportunity to agree to explicit uses of their personal data) features prominently as a key cornerstone of privacy and data policy.

The [2017-2018 Annual Report from the Office of the Privacy Commissioner \(OPC\)](#) highlighted that the office seeks to take an active role with respect to AI being compliant with privacy laws. The OPC has expressed that consent may not always be practicable in the context of AI, where such vast amounts of data are being collected and used for different purposes. Other forms of protection are necessary in such cases, and are actively being investigated. Effective AI governance will enable this investigation by allowing the department to make these decisions through a case and context sensitive approach rather than one size fits all cautionary approach.

An additional OPC guidance piece ([Inappropriate Data Practices](#)), prohibits the following data uses:

- Profiling that leads to human rights law violations
- Uses that are likely to cause significant harm to individuals
- Posting personal information with the intent of charging affected individuals for its removal

These use cases represent downright unethical practices that are extremely distant from the standards to which ESDC holds itself, but are worth noting to highlight the OPC's current concerns with respect to data use.

In October 2018, the International Conference of Data Protection and Privacy Commissioners provided the [Declaration on Ethics and Data Protection in Artificial Intelligence](#). This declaration affirms that the respect of the rights to privacy and data protection are increasingly challenged by the development of artificial intelligence and that this development should be complemented by ethical and human rights considerations.

Much of the privacy legislation and policy that a strong AI strategy will need to encompass are still in the process of being created or reformed around the world. Although the experience that ESDC has gained through its pilot work will undoubtedly prove invaluable, it will be imperative that we keep up to date on the ethical, legal and policy conversations related to privacy for sustaining responsible AI development and implementation over time.

Considerations and Risks

Many issues related to data and privacy are not AI specific. AI Policy will accordingly need to refer heavily to the broader ESDC Data Strategy and Data Policy on many overlapping issues, and not deviate from the tone set by these pieces.

At some point, with near certainty, a particular AI initiative will face a trade-off between model effectiveness (and increased service/value for the taxpayer) and citizen privacy (via data collection and storage). It will be critical that this trade-off be explored in its entirety and public opinion solicited through consultation and other means. Modern AI at ESDC will be constrained if its related policy and practices are not reflective of modern public views on these issues. This is why ESDC will need to a case- and context-sensitive risk-management approach to identify, monitor and mitigate AI privacy risks.

Current Status and Plan

The AI governance proposal embeds privacy considerations directly into the process of developing AI solutions regardless who or where in the department they are being created. Via the Data and Privacy Committee (DPC), the department has organized itself to have data more prominently featured in privacy discussion, enabling a more direct conduit to address data- driven AI and privacy issues. AI projects will be presented to DPC, where privacy experts including the Chief Privacy Officer and from the Privacy Management Division will be relied upon to consider the privacy risks of a given project.

Additionally, the Chief Data Office will also actively be investigating and testing privacy-preserving data science options as they become relevant (such as federated learning and homomorphic encryption). These initiatives will allow the department to find attractive points in the utility/privacy trade-off space.

4.3c Transparency

ESDC has several responsibilities to the public and other stakeholders with respect to transparency:

- We need to adequately explain decisions that affect clients
- Our business processes need to be subject to independent review/audit to ensure the department is acting responsibly
- ESDC, as part of the federal government family, has expressed an increased desire to be open by default for many of its activities, being mindful of integrity concerns.

These responsibilities will continue to be present as we integrate AI into our systems and operations, and though their overall spirit will not change, the specifics with respect to their implementation will as we develop new approaches.

Further, like almost no field of study before it, artificial intelligence is remarkably open by nature, and this openness is driven by the leaders in the field (Google, Microsoft, Apple, FaceBook, Amazon and IBM). ESDC will need to find its place in this community, noting that it's not making investments at the level of the big players, but still can be a valuable contributor.

Current Requirements and Policy

The call for AI models to be interpretable is coming from every area of the department, GoC, academia and the public at large. The TBS [Directive on Automated Decision Making](#) requires that a "meaningful explanation" be provided to affected individuals of how and why a decision was made. Currently, this policy remains at a high level, providing limited description of what is meant by a "meaningful" decision.

Treasury Board has also stressed a number of initiatives related to [Government, including the Open First white paper and Digital Playbook](#). These initiatives represent a new way of thinking for the Government of Canada in an effort to be transparent.

ESDC's current policy is that it does not make open its business processes, for reasons mainly associated with integrity risk.

Considerations and Risks

Modern AI models are often "non-interpretable", meaning they are not able to explain the factors that led to specific outputs as have been traditional statistical models (this feature is effectively traded away for powerful problem solving abilities and improved model accuracy).

In cases for which the implications of the decision pose limited risk, this lack of explainability is acceptable (e.g. a model that advances a file in the queue for efficiency reasons). In other cases, the department has a direct duty to explain administrative decisions, and this will not change with the introduction of AI. However, excluding algorithms because they are not readily explainable could decrease the usefulness of an AI

tool. Well thought-out policy pertaining to the level of explainability of models designed for different purposes will be needed accordingly.

Luckily, model explainability is a very active area of research. For example, developing research in [XAI](#) aims to decipher the reasons for a decision made by AI, or to create new machine learning methods that are transparent. We are monitoring developments in this area and will investigate how they can be incorporated into our systems.

One intriguing aspect of recent AI research, methodology and use-cases is that *low risk applications* have been made remarkably open and accessible to anyone who is interested. This serves several purposes:

- They make for readily accessible and effective training solutions for AI practitioners.
- Scrutiny can be applied by the open source community to a degree that is effectively impossible in closed environments. Vulnerabilities have the potential to be detected and patched long before being exploited.
- Models, training data and code can be leveraged for other organizations that may directly or indirectly provide benefits to both the department and its clients.

There are, however, potential complications with the full openness of AI methods in our context:

- Fraud and program security
- Given that models will fit into the broader business processes, it may be difficult to make the AI components open given that business processes aren't currently published.

Current Status and Plan

The Chief Data Office has significant experience with the inner workings of AI models, is expected to play a critical role in aspects of AI Policy that pertain to model interpretability. Business areas will play a central role related to the level of explainability required for respective programs. It is desired that AI Policy get very specific with respect to which AI use cases require different acceptable levels of model explainability.

On the open front, integrity concerns are expected to trump the sentiment for openness with respect to high risk business processes (administrative decisions with significant financial implications being the highest risk). The Transformation and Integrated Service Management Branch will play a critical role with respect to AI transparency policy for administrative decisions. Transparency pertaining to processes with less inherent risk will be a broader discussion.

4.3d Integrity & Security

As stewards of approximately \$130 billion dollars of annual benefit payments and the personal data of every Canadian with a Social Insurance Number, the department has every obligation to secure its systems and data. This represents an especially challenging task in the world of artificial intelligence, open source algorithms, API calls and cloud computing.

Considerations and Risks

Although risks associated with ESDC's business processes being compromised have been present since their inception, AI presents new dimensions to these risks. Some potential security threats and issues affecting the integrity of AI and machine learning models are the following:

- **Reverse-Engineering [1A]:** Computer scientists from Cornell Tech, Swiss Institute EPFL, and the University of North Carolina replicated the output of Amazon AI and BigML models by analyzing a few thousand query-response pairs. The dangers of reverse-engineering are that a black-box can become functionally known and leveraged by attackers, or attackers can implement the stolen model in their own intellectual property without the owner's consent.
- **Adversarial Injection Attack [2A, 2B, 2C]:** This attack learns how to change the input of a model slightly with the goal of disrupting it and triggering misclassification. This attack can be carried out in such a way that a change to the input is not noticeable via manual inspection, but causes the model to predict a complete different outcome. For example, it's possible by changing just a few of the right pixels in an image of a Quebec driver's license, it will cause the model to predict the presence of an Ontario driver's license card. In order to carry out an injection attack, the attacker requires a reverse-engineered (or the original) model. By using the reverse engineered model, the attacker can algorithmically determine the minimal amount of injection required alter the models outputs. The following are some ways in which an attacker can obtain the information needed to produce successful adversarial injection attacks:
 - **Model Parameter Leak:** If the black-box model is weights or parameters are leaked (e.g., the weights and architecture of speech recognizing deep neural network), predictions can be generated through exploration of their outputs. In the case of an understandable white-box model's parameters leaking (e.g., a decision tree rules), the attacker will have direct access to the models decision making.
 - **Training Data Leaks:** If training data is leaked, an attacker could leverage from this data sensitive information revealing vulnerabilities in the machine learning model.
 - **Exploratory Attacks Through Indirect Model Access:** This is when the attacker will try to understand how the model predicts by testing various inputs and measuring outputs. (e.g., trying various inputs in a web portal and observing the results)

- **Adversarial Dataset Attacks:** The attacker will use adversarial injection attacks to send adversarial data to the system, the goal of which is to have the adversarial data included in the new training data. The training dataset becomes an "adversarial dataset", which if used can corrupt a Machine Learning model. This type of attack is mostly targeted towards models that receive regular feedback. For example, a model that improves a websites search results based on user feedback can be spammed with misleading feedback, the faulty feedback could then be added to the models training data and cause a disruption of the search service. Adversarial dataset attacks can affect both false positives (e.g., deny services to others) and false negatives (e.g., to gain illegitimate access to services).

Current Status and Plan

As part of the AI governance design, the implementation phase includes explicit considerations on how particular AI solutions will be monitored and reviewed over time, including how they will be protected from third party access and reverse engineering. The Departmental Security Officer and IITB Security lead the security function for automated decision systems in the department, and will play a critical role as security threats take on new forms in the AI environment.

4.3e Bias

Bias is a concern in the development of both policy and human-performed tasks, potentially stemming from cognitive bias or social beliefs, and AI systems conceptually have the ability to overcome these obstacles. However, bias in AI tasks (e.g. decision-making, classification) can also occur depending on the data or algorithm used to create a predictive model, and can have significant consequences when a machine is performing tasks much faster and on a larger scale than a human could. AI bias can come in many forms: data set bias, algorithmic bias, systemic prejudicial bias, and procedural bias, to name a few. It will be critical that ESDC AI solutions exhibit the highest standards of procedural fairness, and consequently minimize undesirable bias to the extent possible.

Bias in training data

Supervised machine learning requires labelled training data, for instance, records of human-made classifications or decisions related to a set of variables for each (e.g. grant/denial based on an application). If the human-created labels are biased, then the machine can also learn to make classifications with the same bias.

- E.g. Predictive Policing, like PredPol and COMPAS, use historical policing data to train a model for prediction of crime hotspots and likely perpetrators. Due to bias in that data, populations who are historically over-policed or targeted along income and racial lines are disproportionately identified as crime risks by the AI.

Incomplete, skewed or non-representative datasets

In either supervised or unsupervised methods, if the distribution of categories in the dataset does not adequately reflect the actual distribution, or key variable values are

excluded, then the machine is making predictions on incomplete information and not identifying reliable patterns.

- E.g. Datasets used to benchmark the performance of facial recognition have over-represented male and light-skinned faces by a large majority, thus masking issues in the accuracy of identifying female and darker-skinned faces.

Emergent/Similarity bias

Sometimes bias is intentionally encoded into an algorithm, where the output suggested to a user is meant to be similar to previous searches or personalized based on what the system knows about the user.

- E.g. Facebook News Feed has presented content to users based on what they have already seen, and typically excluded links that take an opposing perspective or on unseen topics, i.e. creating a "bubble", or perpetuating confirmation bias.

The quality of AI, and avoiding bias, require designing the right framework from the beginning, not just checking outputs of developed software. AI Policy should provide sound guidance in this area.

Considerations and Risk

Gaps or historical biases in datasets can cause AI systems to unfairly withhold services, opportunities or resources, which is known as allocative harm. They can also reproduce and amplify harmful stereotypes, causing representational harm. Although technological solutions to reduce biases exist, they are often limited in their capacity to address historical and systemic inequalities. Analyzing datasets for potential biases - and addressing such elements - prior to feeding them into algorithms can help limit unintended harm to individuals and organizations. Some researchers and activists also argue that citizens affected by AI-generated decisions should have the right to see the data, know how it was generated, be able to correct it when necessary and be able to contest decisions.

GBA+ is an analytical process used to assess how diverse groups of women, men and non-binary people may experience policies, programs and initiatives. The "plus" in GBA+ acknowledges that GBA goes beyond biological (sex) and socio-cultural (gender) differences and considers other identity factors, like race, ethnicity, religion, age, and mental or physical disability. As is the case for policies, programs and projects within the Department, the integration of GBA+ in the design, implementation, monitoring and evaluation of AI solutions can help identify and reduce inequalities and bias. For instance, applying GBA+ to ESDC's AI solutions could help identify if certain groups of Canadians are over or under-represented in databases, and, if so, any underlying reasons, and what the potential solutions are.

Another tool, Social Systems Analysis can also be useful to help identify the impacts of AI systems on all parties. Similar to GBA+, it considers the historical, social, political and economic context in which a set of data were produced, including the classification and coding of data. It also examines the extent to which differences in communities' access to information, wealth and basic services shape the data that the AI model is trained on.

Current Status and Plan

The CDO is expected to take the lead in ensuring model biases are properly measured, understood and addressed from a technical and mathematical perspective project-by-project. Business areas, with other AI stakeholders, will decide what types and magnitudes of bias are acceptable for their business processes.

4.3f Impact on the Workforce

The department highly values its world-class staff, and our current and future efforts into the AI world are to augment and relieve employees, not to replace them. The changing nature of work and the advancement of technology are inevitable, so we want to use AI together with staff in inventive ways that enable service delivery options that were not previously possible.

Considerations and Risks

To ensure AI activities align with the objective that current and future AI solutions are to augment employees work and not replace them, AI Policy will need to address specific aspects related to that objective:

- AI automation can take over specific tasks, while boosting productivity and, potentially, demand for a service.
- Changes to employee tasks may result in changes to employee work descriptions and trigger changes to classifications and organizational design.
- In time, the public will dictate how best to use AI in its public service, and it is our job to provide the public with options/flexibility.
- The growth in AI is opening up new opportunities in emerging technology leading to new job creation and the need for employees to learn new skills.

Current Status and Plan

Currently, no AI solutions being developed in the department present a risk of affecting the workforce to the degree of needing organizational change. As part of AI governance, workforce implications will need to be considered as projects are developing, bringing in relevant stakeholders. Stakeholders from the Human Resources Services Branch, labour relations, unions and the Public Service Commission should consider the impacts a particular solution may have on the number of employees needed in a particular process and changes to the nature of their work.

The ESDC Data Strategy emphasizes the need for investing in people and provides training and career paths for all analytics for all analytics personas from policy and business analysts to data scientists.

4.3g Performance Measurement

As discussed in Section 3, AI needs to be trained with data to do its job. The training data sets provided to it would establish its sense of right and wrong, which would be defined by the business area and respect the governments polices, legislation and ethical standards. From a technical perspective, the department needs to know that the AI solution is following its training and is performing adequately for its function. Furthermore, AI and machine learning models are not static technologies, so their performance will need to be monitored over time.

Considerations and Risks

At ESDC, human decision-making is recorded and managed through a structure and process of delegation of authority and is established through legislation such as the [Financial Administration Act](#). These human decisions are assessed through quality assurance and auditing mechanisms to determine if programs, services and internal operations are following roles and responsibilities appropriately. Decisions made using information from an AI solution or made by a solution will also need to be recorded and managed, but this structure of audibility needs to be hardwired into the AI solution or documented by users. Regardless if decisions are made by a human or technology, the department evaluates the outcomes of these decisions to determine if objectives are achieved.

Depending on the business process in which the AI solutions is deployed, the department will need to decide:

- How we will know that an AI solution is making a "right" decision and providing the adequate outputs.
- How the business area will know whether an AI solution is applying the same rigour of analysis as a trained, experienced human agent.
- How often the solution should be audited and if audits should they be transparent to the public.
- How performance and outcomes will be monitored over time.

Current Status and Plan

The government's current focus on Results and Delivery has required the department to reexamine its business models and processes with a view of achieving better results for Canadians. A key part of this shift is the establishment of timely, complete, accurate and relevant performance information to inform decision makers about programs and services. When investing in AI solutions this same performance examination is needed.

When developing AI solutions key performance indicators (KPIs) should be established to provide a benchmark for success. We cannot use AI just because it is cool. It needs to add long-term value for the citizen. As part of the AI governance design, during the implementation phase the business area will establish KPIs, including ongoing monitoring of data bias. As part of ongoing review and monitoring, ESDC's Evaluation function will not look at an AI solution in and itself but rather the outcomes achieved by its use; while

the process of the decision-making will be what ESDC's Internal Audit will provide assurance around.

4.3h Value

Delivering value to department should always be top of mind during all of our business transformation and service improvement initiatives. Proposals to implement an AI solution should always aim to help reduce costs, enhance program integrity, achieve better performance and results, and improve service delivery to Canadians.

Section 5 of this strategy provides an in-depth discussion on the AI value proposition illustrating the importance of valuing our data, building up internal capacity through a robust data analytics program, investing in our people and obtaining maximum value from our vendor arrangements. The AI governance design expresses the need to consider what value a project will bring by allowing the business area to research and explore while bringing in relevant stakeholders to assess return on investment.

4.3i Other Policy Considerations

Open Source Software

Open Source Software (OSS) is commonly used when developing AI solutions. OSS provides value to the department because it can generate an increasingly more diverse scope of design perspective by leveraging knowledge throughout the industry. However, even when freely available, OSS is governed by licensing conditions, which imply strict contractual restrictions. Lack of awareness regarding the existence of such restrictions poses legal risks to the department, both when used in-house and when working with vendors who use OSS in their software development projects.

Some OSS licensing conditions include simple obligations, while other restrictions included in OSS licenses are more complicated and may not provide the intellectual property expectations the department needs. Some conditions include:

- An Infection effect: certain uses of specific OSS can cause the entire resulting software development to be governed by the respective OSS license. The software will be subject to compliance with the very same requirements that were applicable to the OSS component used, turning the resulting solution entirely open.
- Disclosure of complete source code: some OSS includes the obligation to disclose the entire non-OSS source code.
- Commercial distribution not permitted: many OSS licensing policies do not allow users to commercially distribute any deliverable, which includes them. A vendor should not be selling or renting deliverables containing OSS that falls under these licenses.
- License prohibitions: some OSS includes clear license prohibitions to modify or to embed OSS for or within a software deliverable.

To mitigate these risks it is first important for the department to be aware that not all OSS is made equal. It is important for technical staff and management to work together on desired outcomes and conduct an assessment with legal services on the impact of each OSS licensing conditions, including what trade-offs may need to be made to achieve objectives.

Other Ethical Considerations

The considerations discussed above are all important and ethics run through them all. However, an AI solution may be legal, it may not break any privacy principles; it may not pose a high level of bias risk, but it still may not be ethical. Although most technology is designed with the best intentions, it can be difficult to anticipate long-term impacts of a product once it is released and reaches scale.

Effective AI governance and policy can help makers of the technology, project managers, business area experts, engineers, and others get out in front of problems before they happen. AI governance aims to facilitate better product development, faster deployment, and innovation that is more impactful. All while striving to minimize technical and reputational risks.

The department will need to consider the ethical implications of releasing and scaling an AI solution. There are multiple points in the governance design when this should be considered. During the initial governance review, a common sense check should take place as to whether we should proceed or not. The research and exploration phase can also determine that the data and technology will not provide an adequate amount predictability given the decisions we hope to achieve, making it unethical to proceed. In addition, during the design phase we may notice that we will not be able to implement adequate controls to stop the AI solution from causing harm when scaled.

As we develop AI, we should review each project for future scenarios by considering:

- How different users might be affected differently?
- What actions we will take to safeguard privacy, truth in decision-making, democracy, mental health, civic discourse, equality of opportunity, economic stability, or public safety?
- What could we be doing now to get ready for this risky future? Are there any new categories of risk we should pay special attention to now?
- What design, team, or business model choices can actively safeguard users, communities, society, and the organization from future risk?

Recently, a high-level expert group on AI in the European Commission presented their [ethics guidelines for trustworthy artificial intelligence](#). According to the guidelines, trustworthy AI should be: (1) lawful - respecting all applicable laws and regulations; (2) ethical - respecting ethical principles and values; and (3) robust - both from a technical perspective while taking into account its social environment. These guidelines are well suited to help frame our thinking and are in line with our current initiatives.

ESDC needs to align itself with the growing number of voices to ensure that AI is developed in a responsible manner so that we do not lose control. It should be understood, however, that we are (for the foreseeable future) working in the realm of Narrow AI, and humanity has not yet advanced to Artificial General Intelligence.

5. The Artificial Intelligence Value Proposition

The AI value proposition is the intention that innovations such as machine learning will improve services to Canadians by speeding up responses to inquiries, benefit delivery and enable greater insight by automating time-consuming internal processes. This will allow more time for reflection, planning and cost-benefit analysis resulting in better policy development.

Achieving this value proposition requires the department to not only invest, but also reinvent our approach to people, processes and technology. Our AI pilot projects have provided us a grounded understanding on how to generate benefits and mitigate risks in order to maximize outputs. This section sets out a vision for how this value can be achieved by considering what will need to be done to shift from a culture of hierarchical and mechanized development to more modular technical design processes. This shift will affect how we build AI solutions in house; how we procure solutions from vendors; and what technology, training and skills are needed both now and in the future. The opportunity for us to set the stage for this shift is ripe because developments in AI are still early.

5.2 Valuing ESDC's Data

ESDC's greatest asset when it comes to AI is that the department is a prime generator and user of data. Across the department data generation, management, storage and analysis are fundamental tasks, used to provide benefits and services to individuals and business, and to detect non-compliance, evasion and fraud. However, the department vast data assets have been underutilized in the past.

In order to properly evaluate the investments that ESDC makes in artificial intelligence, the department must first properly assess the value it places on its data and how it perceives them. The desire to change how data is being used at ESDC has led the department to create the position of Chief Data Officer and to establish the CDO office. The mandate of the CDO is to maximize the value of ESDC's data assets, to maximize the way data is collected, stored and analyzed.

- **Collection:** Given a particular problem, what data is needed and what is available? Are their limits to what we are legally authorized to do with the data we do have?
- **Data flows and infrastructure:** How do particular types of data flow through the department and is it reliable? Where is structured and unstructured data stored and who has access?

- **Exploring and transforming:** Given a particular problem or context, does the data need to be cleaned? Do we have an incomplete data set? Are there other data sets in the department that can be linked together for a more complete picture?
- **Data analytics:** What are our data stories? Can we use our data to define metrics to track change and our understanding of the data given various factors and contexts? Do we know what we want to predict or learn? Can we create training data by generating labels?
- **Learn and optimize:** Putting in place AI experiments and pilots to be grounded in experience so that we can learn incrementally, to minimize risk and optimize results.

To achieve the CDO mandate, one of the first tasks of the OCDO was to develop the first ESDC data strategy.

Data is a valuable thing and will last longer than the systems themselves. "- Tim Berners-Lee, inventor of the World Wide Web.

In other words, take care of our data, we never know what problems they can help us solve!

ESDC's Data Strategy

The vision of the [ESDC data strategy](#) is to ensure that employees have access to the data when they need it. To achieve this vision, the BDPD has a number of goals to achieve:

- Make data a business asset through effective governance and stewardship
- Make data accessible and secure
- Transform business group and IT group partnerships to improve data use
- Provide people with knowledge and tools to use data
- Make data science, including the ability to develop artificial intelligences, a core competency of the department.

To achieve these goals, the strategy is based on six pillars: data governance, data access, empowerment, people, data management and data science.

How will these pillars enable us to put in place an artificial intelligence strategy?

Data Governance: Governance is about answering the following questions: what data do we have; what do we need; where can we find these data; what is their degree of reliability; Who makes the decisions; which rules should be applied?

Data Management: Data management is about ensuring the infrastructure is in place to securely store data and to provide users with access to embedded data with the tools they need to analyze it.

For purposes of artificial intelligence, data governance and data management are aimed at ensuring that the data is ready, willing and able to be used once the initiatives are ready for launch. It would be unacceptable for the Department to use faulty data and inappropriate processes.

ESDC AI Strategy

DRAFT - V3.0

Access to data: Access is about making data available to all who need it quickly and securely while protecting confidentiality

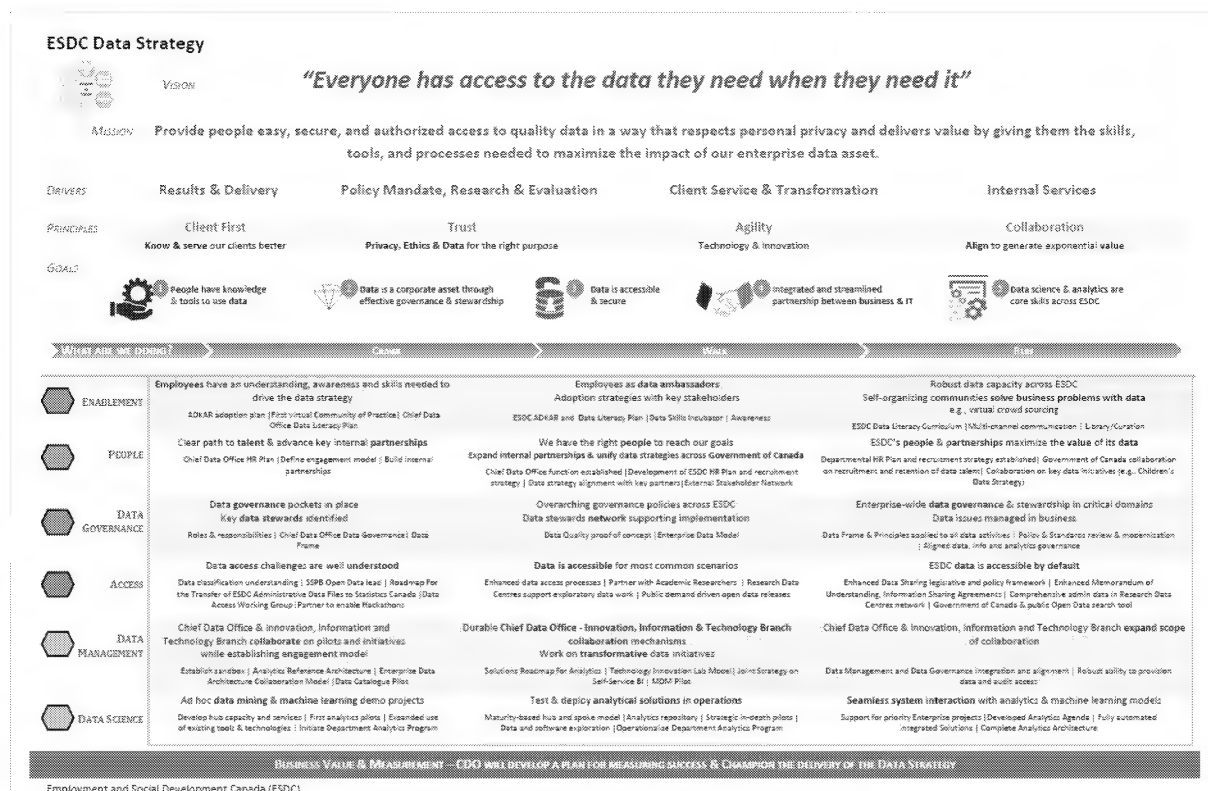
Empowerment: Empowerment is empowering people to make better use of data (understanding of data, cultural change, support, communities, communication and tools)

The pillars of "data access" and "empowerment" are aimed at creating an ESDC culture that promotes innovation and experimentation. Such a culture is paramount in all organizations with great aspirations in the field of artificial intelligence.

People: The People Component is about recruiting and retaining people with the skills we need, creating the right team structures, and working with ESDC partners.

Data Science: Data science involves developing a program to build the analytical capacity to use methods such as machine learning, AI and other methods and tools to discover new information from data analysis.

Data science is the culmination of all the other pillars of the data strategy. In fact, if the data is ready and of good quality, accessible, with the right people present and ready to use it, data science can be effective and can bring the most value to the Department.



The artificial intelligence strategy is an important component of the ESDC analytics program and its value is enabled by a sound data strategy. It defines what people can expect from analytics in the department, the importance of having people with the right knowledge, how to approach the different companies selling artificial intelligence services and infrastructure necessary to optimize the use of public funds.

As can be seen, the ESDC data strategy is dependent on the analytics program while the latter depends on the artificial intelligence strategy. However, it is the latter that benefits most from the data strategy because with the latter in place, all the initiatives presented in the strategy on artificial intelligence can take off.

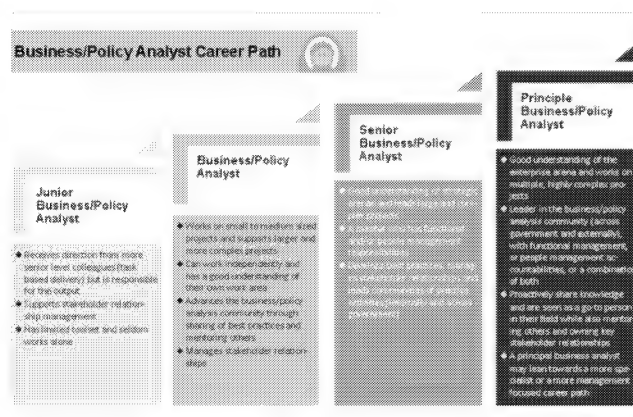
5.3 ESDC's Analytics Program: Building for the Future

Robust data analytics is the corner stone of our ability to provide value through AI. This stream of the [Data Strategy](#) is for the department to build upon its people. This includes the need to focus on recruitment and retention of people with the needed skills; putting in place the right team structures; providing the necessary tools; and keeping up to date on industry standards. In the area of Analytics (within which AI resides), this will be realized by the [ESDC Analytics Program](#).

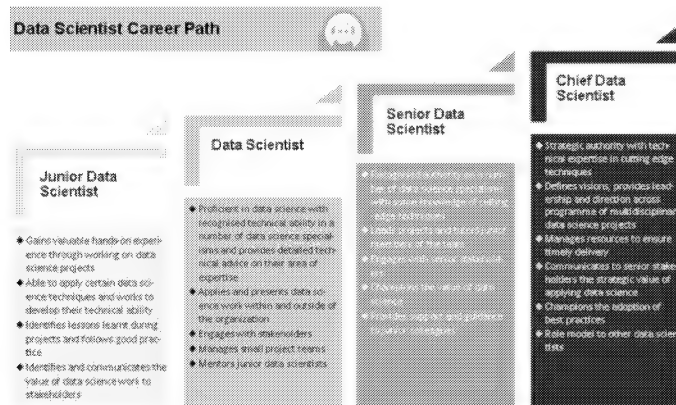
The objective of the Analytics Program is to create a scalable, secure analytics ecosystem that supports all business uses in the department to improve outcomes for citizens, families, organizations and communities. Not all problems that require data analytics need AI to provide useful insights. This is why it is important for the department to have tailored analytics training, awareness and career paths that sustain the right mix of techniques to be applied to different sets of use cases. We must know when AI is the solution to a problem and when it is not. (so we don't ask for a rocket ship, when all we need is a skateboard).

The analytics program establishes training and career paths for all analytics personas. This allows the department to understand and apply the needed roles and responsibilities of all employees involved in the development work and what value each bring to the design process. It enables employees to know what training they need for what role and how it will be beneficial, creating a culture of data literacy continuous learning. In the future this understanding will support the development of workforce adjustments and organizational design shifts brought about by increased automation.

Business / Policy Analyst Career Path



Data Scientist Career Path



Data Literacy Skills and Learning Objectives


Personas	Data Science			Data Governance		Business Analytics	
	Citizen Data Scientist*	Data Scientist	Data Provisioner**	Data Governor		Analytics Consumer*	Business/Policy Analyst
Learning Objectives	<ul style="list-style-type: none"> ● Develop a depth of knowledge and understanding of principles, tools, and applications of data science ● Apply quantitative modelling and data analysis techniques on economic, social, or service delivery issues that are relevant to ESDC as well as determining whether such as Python, R, and SQL ● Gain knowledge of advanced quantitative methods such as data mining, predictive modelling, or machine learning, and artificial intelligence to conduct data analysis of business problems ● Apply a variety of data management techniques to optimize business value of data at various points in the data life cycle, from collection, governance in data sources, and integration ● Effectively communicate findings and present results from data analysis using data visualization techniques 			<ul style="list-style-type: none"> ● Develop the knowledge and understanding of data governance frameworks and processes, and the practices and the disciplines for managing data and information ● Be able to define and understand the roles and responsibilities associated with data stewardship and other roles within the data governance framework, and to ensure data accountability ● Develop and implement from scratch to assess the quality of departmental data and ensure data to data data integrity ● Own the necessary technical knowledge and skills to build, implement, and improve data security safeguards, build quality data data architecture and data, and implement data protection standards ● Define, maintain, and report on a means to evaluate the effectiveness of the data governance program and data stewardship efforts 		<ul style="list-style-type: none"> ● Develop the knowledge of business analysis, project delivery practices, and standards across the project life cycle ● Understand the role and importance of the business/policy analyst ● Develop a solid understanding of the various tasks that require business/policy analysis ● Learn how to plan, elicit, analyze, model, validate, document, and manage user requirements through the project life cycle ● Conduct business/policy analysis and research, provide strategic advice and recommendations, make decisions and assist in management ● Be able to ask the right analysis questions, understand and interpret the findings, and apply analytical insights into business operations ● Effectively communicate findings and present results from business/policy analysis to stakeholders 	
Skills	Technical <ul style="list-style-type: none"> □ Mathematics and statistics □ Programming & database skills (Python, R, SQL, etc.) □ Advanced data analysis & artificial intelligence, or machine learning, text mining, and natural language processing, etc. □ Data management □ Data visualization and presentation Non-Technical <ul style="list-style-type: none"> □ Business domain knowledge □ Business risk assessment and quantitative understanding and dealing with a business situation to deliver better business results □ Problem solving □ Critical thinking 			Technical <ul style="list-style-type: none"> □ Data analysis and reporting □ Data management □ Data quality assurance Non-Technical <ul style="list-style-type: none"> □ Data governance knowledge □ Data modelling □ Data architecture □ Policy development and management 		Technical <ul style="list-style-type: none"> □ Requirement elicitation and formalization □ Requirement analysis and modeling □ Business analysis tools and techniques □ Policy research and synthesis □ Quantitative and qualitative data analysis Non-Technical <ul style="list-style-type: none"> □ Strategic thinking, problem-solving, innovation, and collaboration □ Leadership □ Communication (writing, oral presentation, data visualization, storytelling, etc.) □ Stakeholder management 	
Curriculum	<ul style="list-style-type: none"> ● Data Science Fundamentals ● Introduction to Programming ● IT/Computer Science Model Development ● Statistics and Probability ● Introduction to Tools - R / Python ● Data preparation ● Exploratory Data Analysis ● Advanced Data Analysis <ul style="list-style-type: none"> ○ Machine Learning ○ Data Mining ○ Artificial Intelligence ○ Big Data ● Contemporary Critical Thinking, and Problem Solving ● Data Visualization and Communication ● Data Privacy and Security 			<ul style="list-style-type: none"> ● Data Governance Fundamentals ● Data Governance for Data Stewards ● Data Governance for Business Leaders ● Data Quality Assessment ● Data Analytics and Modeling ● Data Security ● Privacy Act and Regulations ● Information Management Fundamentals ● Data Governance Tools and Techniques ● Governance, Critical Thinking, and Problem Solving ● Data Visualization and Communication 		<ul style="list-style-type: none"> ● Statistics and Business Analysis Fundamentals ● Business Analysis Planning and Monitoring ● Business Analysis Tools & Techniques ● Planning, Eliciting, Developing, Validating, and Managing User Requirements ● Quantitative and Qualitative Policy Analysis ● Policy Writing & Briefing ● Preparing for the IRBA/CCBA (Certification of Competency in Business Analysis) CBAP/Certified Business Analyst Professional Certificate Program ● Creativity, Critical Thinking, and Problem Solving ● Data Visualization and Communication 	

Analytics Program Oversight

There are a wide variety of analytical stakeholders and projects across the organization that has resulted in a patchwork of capabilities across the Department. The CDO will add value by playing an oversight and leadership role, with IT and Business, to manage, mature and optimize the analytics capability in the Department.

Examples of Support Services

- Help you build an analytics team
- Help you find skilled analysts
- Help expand your analytics toolbox of techniques
- Help you connect to Department tools & data
- Help you access training material and training options for tools
- Help you develop/procure new tools & enhance existing solutions
- Help you get started on first project
- Design & complete projects
- Networking & collaboration opportunities



Range of Need:

Start:

I'm starting from scratch but want to build capacity to be self reliant.

Grow:

I have base expertise but want to grow.

One time:

I need help but not feasible to build analytics capacity.

Community of Practice

The mature state of the Analytics Program envisions a strong, department-wide community of practice for analytics. This includes data scientists of course, but also identifies a number of other key players (i.e. analytics consumers, business analysts, data provisioners, apps developers). This environment will foster collaboration, knowledge and resources sharing, open practices, and provide a number of other benefits. The more employees that understand AI and related digital technologies the more diverse the voices, eyes and ears can be present to flag potential risks and impacts.

5.4 The Need for Internal AI Capacity at ESDC

Artificial Intelligence and related digital technologies are unique from traditional technical solutions because much of AI advancement has been because of collaborative development between multiple independent contributors. Open-source software has generated an increasingly more diverse scope of design perspective than any one organization is capable of developing and sustaining long term. This presents an opportunity for the department because to achieve true value in AI development it is critical that we have internal capacity that continues to learn and adapt to changes in the industry.

The need for internal AI capacity was one of the critical lessons learned during our early AI pilot projects. Internal capacity provides:

- The ability to develop custom solutions to save staff time and enable new ways of doing things.
- What we learn, we can leverage from one project to the next. New ideas open up as internal staff skill-up and become more familiar with our business processes.
- We're able to properly evaluate proposed AI work by external vendors (knowledge is power). This enables us to make the right business decision on behalf of Canadians.
- Enables us to demystify AI for the department, enabling more effective discussions and pragmatic AI solutions.

We know we can't build it all, especially while AI remains a relatively young industry as far as private sector solutions go. Also given the nature of Open-source software development, value can only be realized for Canadians through ESDC having its own internal capacity in AI. The department will proactively train and educate its workforce; recruit top talent; and engage with private enterprise, academia, citizens, as well as other government departments and jurisdictions.

ESDC's Machine Learning Seminar

Since December 2016, the Data Science Division of the Chief Data Office has been hosting a weekly machine learning seminar at the ESDC Learning Centre.

Year 1 featured lecture-style presentations that introduced concepts like machine learning paradigms, deep learning and their many applications to the department.

Year 2 has been more projects focused, where seminar participants actively work together to solve AI problems:

- Several text classification solutions that identify when toxic comments are present in an internet forum
- Developed an automated scraping/analysis tool for the Canada Public Servants subreddit

- Actively working on an information retrieval project for the canada.ca website
- Methods and applications of reinforcement learning

In addition to ESDC staff, we've had attendance from several other organizations as well, including Statistics Canada and the Department of Justice. The seminar continues to evolve, but enjoys an enthusiastic, hard-working environment that is driven by its energetic membership. All ESDC staff are welcome!

Please check out our [GitLab page](#) for past topics, code repositories and more.

5.5 Obtaining Maximum Value with Vendors

As with any area of public service life, getting maximum value for the taxpayer should be a primary goal. As we are at the dawn of a new era of investment with respect to AI, there is an invaluable opportunity right now to properly set precedents with AI vendors on behalf of the Canadian public. The open source and collaborative nature of AI advancements presents special considerations when procuring these technologies.

Most of government procurement has operated in an environment of static technical requirements and the purchasing of solutions to well-known problems. However, AI and related digital technologies present a different reality, where development often occurs in a modular technical design process. Problems may be known but the project is completed through multiple development cycles known as iterations. Each iteration is reviewed and critiqued by the team where insights are gained and used to determine what the next step should be in the project. The solution then becomes the sum of the iterative cycles.

This reality has critical implications for how the department should consider procurement decisions that impact elements of contract management such as intellectual property, statements of work, goods vs. service determination and more. If a solution is going through multiple iterations, some done internally and others done with a vendor, will we have control over all aspects of such a development process? Particularly, the intellectual property of the tools used or results obtained, to avoid any risk of losing valuable assets such as an algorithm that can be adopted in other projects.

With these considerations in mind, we put forward vendor-guiding principles that will strategically define the manner in which we think about, assess and view our future vendor relationships when procuring AI solutions.

Vendor Guiding Principles

1. ESDC's data is valuable; it should not be considered a by-product of service delivery or program policy. Its value (and access to it) should be managed with vendors accordingly;
2. If we pay the development costs for a solution, we should not be paying perpetual access fees;

3. If we understand how the technology works and how much it would cost to build it in house, we are in a much better contracting position;
4. If we have full access to the underlying algorithms and models, we can leverage them for other purposes;
5. If we're flexible in the design, we can move components around and won't be forced to continue with arrangements that no longer work for us;
6. There are hundreds of AI vendors. There is only one ESDC.

The key of these guiding principles is to provide future flexibility to the public in how it might choose to democratize the benefits of AI. That choice cannot be made for purchases that are closed one-offs, that cannot be leveraged for other projects and for which the department is paying perpetual access fees.

Knowledge is Power

All other initiatives outlined in this strategy drive towards putting ESDC in a better position for its AI purchases:

- An effective communication strategy results in an extremely well-informed organization that knows exactly the worth of what it is buying.
- A strong internal capacity is needed to undertake the work ourselves if the right deal isn't out there, and to verify that vendors are delivering what they promised.
- Proper infrastructure and processes are needed to provide flexibility in solution deployment.

Government of Canada Procurement Processes

As AI, governance matures so will many of our internal processes and standards related to procurement, contracting and IP reflecting the policy considerations discussed in section four. The current procurement structure at ESDC does not effectively support an iterative design process. This issue is not unique to our department, but is something all GoC departments are confronted with.

To address this issue Public Services and Procurement Canada (PSPC), together with the Treasury Board of Canada Secretariat (TBS), established a list of suppliers who can provide the Government of Canada with responsible and effective AI services, solutions and products.

Considerations and Risks

It is a positive development that the government is providing more and more innovative procurement options for departments to take advantage of when developing AI solutions. However, ESDC must approach these pre-qualified arrangements with caution because the contracts may not fit the needs, requirements and legal obligations of the department.

When deciding to use streamlined procurement vehicles, we must keep in mind the above ESDC vendor guiding principles and be aware that when entering these arrangements:

- The business area should still be going through the AI governance process
- ESDC may not be the contracting authority and should not be dealing with the vendor directly in the case of contract issues
- We must consult relevant stakeholders such as the IP Center of Excellence and legal services branch if changes need to be made to the pre-qualified statements of work
- We are still governed by OSS are licensing conditions discussed in 4.2i.

6. Links to Other Departmental Initiatives

Data Strategy and Analytics Program

As the current wave of artificial intelligence is largely data driven, and the techniques are grounded in Data Science, ESDC has also begun to orient itself towards AI through other initiatives. The (link to be added)ESDC Data Strategy, being led out of the Chief Data Office, aims to maximize the impact of our enterprise data asset, through modernizing how the Department approaches data governance, data management, data access, security and privacy.

Further, the (link to be added)Analytics Program, a key component of the Data Strategy, will form and scale up a frame for a burgeoning ecosystem of analytics across ESDC, to exponentially increase the insights we derive from our data. The Analytics Program includes initiatives for enhancing departmental capacity in delivering analytics solutions, enabling our staff with modern infrastructure and technology for analytics endeavors, and putting in place proper processes and oversight for deriving robust, high quality insights.

Service Transformation Plan

The Department developed the Service Transformation Plan (STP) to support its move from strategy to implementation for transformation and modernization of its services as it moves forward in advancing it's vision for improved service delivery to clients. The Service Strategy is the departmental modernization plan of action that will transform the way we deliver service so that, in the future, Canadians will be able to digitally self-serve, access services seamlessly, receive high-quality, timely, accurate services, have their needs anticipated and receive service from a well-equipped, knowledgeable workforce.

The solutions in the Service Transformation Plan are organized in five groupings based on their impact on clients and similar capabilities:

Allow Me Allow citizens and clients to access their services/benefits in a faster and efficient manner.

Trust Me Enable better ability for clients to apply for benefits/services faster by leveraging know data about the client. Clients will feel trusted and recognized.

Tell me Give more information about the benefits and services and have multiple means of efficiently communicating.

Hear Me, Show Me Increased ability for clients to provide feedback and answer their questions.

My Choice Provide multiple options to engage with ESDC so clients have their choice in how they want to interact and receive benefits/ services.

More information on the Service Transformation Plan can be found [here](#).

(Placeholder for a blurb introducing STP, and see what links are available that we can direct readers to for further reading)

Several solutions in STP are actively investigating the use of AI to enable superior service. Some examples include:

- Solution 2.4: Document Upload, which will provide flexibility to clients in how they provide us with information, and is exploring the use of Computer Vision AI to automatically get the needed information into our systems.
- Solution 3.5: Chatbot, which will provide clients with the ability to first interact with a digital agent to resolve issues, speeding up resolution times.
- Solution 4.2: Program Knowledge Repository, which is investigating the use of modern information retrieval techniques (smart search) to retrieve information from our websites, manuals and other information stores in an automated, efficient manner.

Integrated Service Management

At a minimum, can link Russell Egan's blog post about leveraging the power of digital and using assistive technologies such as voice-over technology, text over images, video captioning, all of which are powered by AI.

Putting it all together

The initiatives outlined in this strategy will provide critical support to all of these initiatives that intend to leverage AI, to ensure the department and its clients can realize broad-ranging, long-term value from these investments.

Data Scientist Career Path



Junior Data Scientist

- ◆ Gains valuable hands-on experience through working on data science projects
- ◆ Able to apply certain data science techniques and works to develop their technical ability
- ◆ Identifies lessons learnt during projects and follows good practice
- ◆ Identifies and communicates the value of data science work to stakeholders

Data Scientist

- ◆ Proficient in data science with recognized technical ability in a number of data science specialisms
- ◆ Provides detailed technical advice on their area of expertise
- ◆ Applies and presents data science work within and outside of the organization
- ◆ Engages with stakeholders
- ◆ Manages small project teams
- ◆ Mentors junior data scientists

Senior Data Scientist

- ◆ Recognized authority on a number of data science specialisms with some knowledge of cutting-edge techniques
- ◆ Engages with senior stakeholders
- ◆ Champions the value of data science
- ◆ Provides supports and guidance to junior colleagues

Chief Data Scientist

- ◆ Strategic authority with technical expertise in cutting edge techniques
- ◆ Defines visions, provides leadership and direction across a programme of multidisciplinary data science projects
- ◆ Manages resources to ensure timely delivery
- ◆ Communicates to senior stakeholders the strategic value of applying data science
- ◆ Champions the adoption of best practices
- ◆ Role model to other data scientists

Business/Policy Analyst Career Path



Junior Business/Policy Analyst

- ◆ Receives direction from more senior level colleagues (task-based delivery), but is responsible for the output
- ◆ Supports stakeholder relationship management
- ◆ Has limited toolset and seldom works alone

Business/Policy Analyst

- ◆ Works on small to medium sized projects and supports larger and more complex projects
- ◆ Can work independently and has a good understanding of their own work area
- ◆ Advances the business/policy analysis community through sharing of best practices and mentoring others
- ◆ Manages stakeholder relationships

Senior Business/Policy Analyst

- ◆ Good understanding of strategic arenas and leads large and complex projects
- ◆ A mentor who has functional and/or people management responsibilities
- ◆ Develops best practices, manages stakeholder relationships, and leads communities of practice activities (internally and across government)

Principle Business/Policy Analyst

- ◆ Good understanding of the enterprise arena and works on multiple, highly complex projects
- ◆ Leader in the business/policy analysis community (across government and externally), with functional management or people management accountabilities, or a combination of both
- ◆ Proactively shares knowledge and are seen as a go-to person in their field while also mentoring others and owning key stakeholder relationships
- ◆ A principal business analyst may lean towards a more specialist or a more management-focused career path



VISION

"Everyone has access to the data they need when they need it"

MISSION

Provide people easy, secure, and authorized access to quality data in a way that respects personal privacy and delivers value by giving them the skills, tools, and processes needed to maximize the impact of our enterprise data asset.

DRIVERS

Results & Delivery

Policy Mandate, Research & Evaluation

Client Service & Transformation

Internal Services

PRINCIPLES

Client First

Trust

Agility

Collaboration

Know & serve our clients better

Privacy, Ethics & Data for the right purpose

Technology & Innovation

Align to generate exponential value

GOALS



1 People have knowledge & tools to use data



2 Data is a corporate asset through effective governance & stewardship



3 Data is accessible & secure



4 Integrated and streamlined partnership between business & IT



5 Data science & analytics are core skills across ESDC

WHAT ARE WE DOING?

CRAWL

WALK

RUN



ENABLEMENT

Employees have an understanding, awareness and skills needed to drive the data strategy

ADKAR adoption plan | First virtual Community of Practice | Chief Data Office Data Literacy Plan



PEOPLE

Clear path to talent & advance key internal partnerships

Chief Data Office HR Plan | Define engagement model | Build internal partnerships



DATA GOVERNANCE

Data governance pockets in place
Key data stewards identified

Roles & responsibilities | Chief Data Office Data Governance | Data Frame



ACCESS

Data access challenges are well understood

Data classification understanding | SSPB Open Data lead | Roadmap For the Transfer of ESDC Administrative Data Files to Statistics Canada | Data Access Working Group | Partner to enable Hackathons



DATA MANAGEMENT

Chief Data Office & Innovation, Information and Technology Branch collaborate on pilots and initiatives while establishing engagement model

Establish sandbox | Analytics Reference Architecture | Enterprise Data Architecture Collaboration Model | Data Catalogue Pilot



DATA SCIENCE

Ad hoc data mining & machine learning demo projects

Develop hub capacity and services | First analytics pilots | Expanded use of existing tools & technologies | Initiate Department Analytics Program

Employees as data ambassadors

Adoption strategies with key stakeholders

ESDC ADKAR and Data Literacy Plan | Data Skills Incubator | Awareness

We have the right people to reach our goals

Expand internal **partnerships & unify** data strategies across **Government of Canada**

Chief Data Office function established | Development of ESDC HR Plan and recruitment strategy | Data strategy alignment with key partners | External Stakeholder Network

Overarching governance policies across ESDC
Data stewards network supporting implementation

Data Quality proof of concept | Enterprise Data Model

Data is accessible for most common scenarios

Enhanced data access processes | Partner with Academic Researchers | Research Data Centres support exploratory data work | Public demand driven open data releases

Durable Chief Data Office - Innovation, Information & Technology Branch collaboration mechanisms

Work on transformative data initiatives

Solutions Roadmap for Analytics | Technology Innovation Lab Model | Joint Strategy on Self-Service BI | MDM Pilot

Test & deploy analytical solutions in operations

Maturity-based hub and spoke model | Analytics repository | Strategic in-depth pilots | Data and software exploration | Operationalize Department Analytics Program

Robust data capacity across ESDC

Self-organizing communities solve business problems with data
e.g., virtual crowd sourcing

ESDC Data Literacy Curriculum | Multi-channel communication | Library/Curation

ESDC's people & partnerships maximize the **value** of its data

Departmental HR Plan and recruitment strategy established | Government of Canada collaboration on recruitment and retention of data talent | Collaboration on key data initiatives (e.g., Children's Data Strategy)

Enterprise-wide data governance & stewardship in critical domains
Data issues managed in business

Data Frame & Principles applied to all data activities | Policy & Standards review & modernization | Aligned data, info and analytics governance

ESDC data is accessible by default

Enhanced Data Sharing legislative and policy framework | Enhanced Memorandum of Understanding, Information Sharing Agreements | Comprehensive admin data in Research Data Centres network | Government of Canada & public Open Data search tool

Chief Data Office & Innovation, Information and Technology Branch expand scope of collaboration

Data Management and Data Governance integration and alignment | Robust ability to provision data and audit access

Seamless system interaction with analytics & machine learning models

Support for priority Enterprise projects | Developed Analytics Agenda | Fully automated integrated Solutions | Complete Analytics Architecture

BUSINESS VALUE & MEASUREMENT – CDO WILL DEVELOP A PLAN FOR MEASURING SUCCESS & CHAMPION THE DELIVERY OF THE DATA STRATEGY